arXiv:1510.05679v2 [math.LO] 17 Apr 2016

A New Notion of Cardinality for Countable First Order Theories

Douglas Ulrich, Richard Rast, and Michael C. Laskowski*
Department of Mathematics
University of Maryland, College Park

April 19, 2016

Abstract

We define and investigate HC-forcing invariant formulas of set theory, whose interpretations in the hereditarily countable sets are well behaved under forcing extensions. This leads naturally to a notion of cardinality $||\Phi||$ for sentences Φ of $L_{\omega_1,\omega}$, which counts the number of sentences of $L_{\infty,\omega}$ that, in some forcing extension, become a canonical Scott sentence of a model of Φ. We show this cardinal bounds the complexity of $(\mathrm{Mod}(\Phi), \cong)$, the class of models of Φ with universe ω, by proving that $(\mathrm{Mod}(\Phi), \cong)$ is not Borel reducible to $(\mathrm{Mod}(\Psi), \cong)$ whenever $||\Psi|| < ||\Phi||$. Using these tools, we analyze the complexity of the class of countable models of four complete, first-order theories T for which $(\mathrm{Mod}(T), \cong)$ is properly analytic, yet admit very different behavior. We prove that both 'Binary splitting, refining equivalence relations' and Koerwien's example [11] of an eni-depth 2, ω-stable theory have $(\mathrm{Mod}(T), \cong)$ non-Borel, yet neither is Borel complete. We give a slight modification of Koerwien's example that also is ω-stable, eni-depth 2, but is Borel complete. Additionally, we prove that $I_{\infty,\omega}(\Phi) < \beth_{\omega_1}$ whenever $(\mathrm{Mod}(\Phi), \cong)$ is Borel.

1 Introduction

In their seminal paper [3], Friedman and Stanley define and develop a notion of *Borel reducibility* among classes C of structures with universe ω in a fixed, countable language L that are Borel and invariant under permutations of ω. It is well known (see e.g., [9] or [4]) that such classes are of the form $\mathrm{Mod}(\Phi)$, the set of models of Φ whose universe is precisely ω for some sentence $\Phi \in L_{\omega_1,\omega}$. A *Borel reduction* is a Borel function $f : \mathrm{Mod}(\Phi) \to \mathrm{Mod}(\Psi)$ that satisfies $M \cong N$ if and only if $f(M) \cong f(N)$. One says that $\mathrm{Mod}(\Phi)$ is *Borel reducible* to $\mathrm{Mod}(\Psi)$, written $(\mathrm{Mod}(\Phi), \cong) \leq_B (\mathrm{Mod}(\Psi), \cong)$ or more typically $\mathrm{Mod}(\Phi) \leq_B \mathrm{Mod}(\Psi)$, if there is a Borel reduction $f : \mathrm{Mod}(\Phi) \to \mathrm{Mod}(\Psi)$; and the two classes are *Borel equivalent* if there are Borel reductions in both directions. As Borel reducibility is transitive, $\leq_B$ induces a pre-order on $\{\mathrm{Mod}(\Phi) : \Phi \in L_{\omega_1,\omega}\}$. In [3], Friedman and Stanley show that among Borel invariant classes, there is a maximal class with respect to $\leq_B$. We say Φ is *Borel complete* if it is in this maximal class. Examples include the theories of graphs, linear orders, groups, and fields.

It is easily seen that for any $\Phi \in L_{\omega_1,\omega}$, the isomorphism relation on $\mathrm{Mod}(\Phi)$ is analytic, but in many cases it is actually Borel. The isomorphism relation on any Borel complete class is properly analytic (i.e., not Borel) but prior to this paper there were few examples known of classes $\mathrm{Mod}(\Phi)$ where the isomorphism relation is properly analytic, but not Borel complete. Indeed, the authors are only aware of the example of p-groups that first appeared in [3]. The class of p-groups is expressible by a sentence of $L_{\omega_1,\omega}$, but no first-order example of this phenomenon was known.

To date, the study of classes $(\mathrm{Mod}(\Phi), \cong)$ is much more developed when the isomorphism relation is Borel than when it is not. Here, however, we define the set of *short* sentences Φ (see Definition 3.6), that properly

*All three authors were partially supported by NSF Research Grant DMS-1308546.

contain the set of sentences for which isomorphism is Borel but exclude the Borel complete sentences, and develop criteria for concluding $\mathrm{Mod}(\Phi) \not\leq_B \mathrm{Mod}(\Psi)$ among pairs Φ, Ψ of short sentences. Furthermore, Theorem 3.10 asserts that there can never be a Borel reduction from a non-short Φ into a short Ψ. From this technology, we are able to discern more about $\mathrm{Mod}(\Phi)$ when $\cong$ is Borel. For example, Corollary 3.12 asserts that

If $\cong$ is Borel on $\mathrm{Mod}(\Phi)$, $I_{\infty,\omega}(\Phi) < \beth_{\omega_1}$.

That is, among models of Φ of any cardinality, there are fewer than $\beth_{\omega_1}$ pairwise back-and-forth inequivalent models.

In this paper we work in ZFC, and so formalize all of model theory within ZFC. We do this in any reasonable way. For instance, countable languages L are construed as being elements of HC (the set of hereditarily countable sets) and the set of sentences of $L_{\omega_1,\omega}$ form a subclass of HC. Moreover, basic operations, such as describing the set of subformulas of a given formula, are absolute class functions of HC. As well, if $L \in \mathrm{HC}$ is a language, then for every L-structure M, the canonical Scott sentence $\mathrm{css}(M)$ is in particular a set; and if M is countable then $\mathrm{css}(M) \in L_{\omega_1,\omega}$ is an element of HC. Moreover the class function css is highly absolute.

One of our central ideas is the following. Given $\Phi \in L_{\omega_1,\omega}$, let $\mathrm{CSS}(\Phi) \subseteq \mathrm{HC}$ denote the set $\{\mathrm{css}(M) : M \in Mod(\Phi)\}$. Then any Borel map $f : \mathrm{Mod}(\Phi) \to \mathrm{Mod}(\Psi)$ induces an HC-definable function $f^* : \mathrm{CSS}(\Phi) \to \mathrm{CSS}(\Psi)$. This leads us to the investigation of definable subclasses of HC and definable maps between them. We begin by restricting our notion of classes to those definable by *HC-forcing invariant* formulas (see Definition 2.1). Three straightforward consequences of the Product Forcing Lemma allow us to prove that these classes are well-behaved. It is noteworthy that we do not define HC-forcing invariant formulas syntactically. Whereas it is true that every Σ_1-formula is HC-forcing invariant, determining precisely which classes are HC-forcing invariant depends on our choice of $\underline{\mathrm{V}}$.

The second ingredient of our development is that every set A in $\underline{\mathrm{V}}$ is *potentially in HC*, i.e., there is a forcing extension $\underline{\mathrm{V}}[G]$ of $\underline{\mathrm{V}}$ (indeed a Levy collapse suffices) such that $A \in \mathrm{HC}^{\underline{\mathrm{V}}[G]}$. Given an HC-forcing invariant $\varphi(x)$, we define φ_{ptl} – that is, the *potential solutions to* φ – to be those $A \in \underline{\mathrm{V}}$ for which $(\mathrm{HC}^{\underline{\mathrm{V}}[G]}, \in) \models \varphi(A)$ whenever $A \in \mathrm{HC}^{\underline{\mathrm{V}}[G]}$. The definition of HC-forcing invariance makes this notion well-defined. Thus, given a sentence Φ, one can define $\mathrm{CSS}(\Phi)_{\mathrm{ptl}}$, which should be read as the class of 'potential canonical Scott sentences' i.e., the class of all $\varphi \in L_{\infty,\omega}$ such that in some forcing extension $\underline{\mathrm{V}}[G]$, φ is the canonical Scott sentence of some countable model of Φ. We define Φ to be *short* if $\mathrm{CSS}(\Phi)_{\mathrm{ptl}}$ is a **set** as opposed to a proper class and define the *potential cardinality* of Φ, denoted $||\Phi||$, to be the (usual) cardinality of $\mathrm{CSS}(\Phi)_{\mathrm{ptl}}$ if Φ is short, or ∞ otherwise. By tracing all of this through, with Theorem 3.10(2), we see that

If $||\Psi|| < ||\Phi||$, then there cannot be a Borel reduction $f : \mathrm{Mod}(\Phi) \to \mathrm{Mod}(\Psi)$.

Another issue that is raised by our investigation is a comparison of the class the potential canonical Scott sentences $\mathrm{CSS}(\Phi)_{\mathrm{ptl}}$ with the class $\mathrm{CSS}(\Phi)_{\mathrm{sat}}$, consisting of all sentences of $L_{\infty,\omega}$ that are canonical Scott sentences of some model $M \models \Phi$ with $M \in \underline{\mathrm{V}}$. Clearly, the latter class is contained in the former, and we call Φ *grounded* (see Definition 3.8) if equality holds. We show that the incomplete theory REF of refining equivalence relations is grounded. By contrast, the theory TK, defined in Section 6, is a complete, ω-stable theory for which $|\mathrm{CSS}(\mathrm{TK})_{\mathrm{sat}}| = \beth_2$, while $\mathrm{CSS}(\mathrm{TK})_{\mathrm{ptl}}$ is a proper class.

Sections 4-6 apply this technology. Section 4 discusses continuous actions by compact groups on Polish spaces. In addition to being of interest in its own right, the results there are also used in Section 6. Sections 5 and 6 discuss four complete first-order theories that are not very complicated stability-theoretically, yet the isomorphism relation is properly analytic in each case. We summarize our findings by:

REF(inf) is the theory of 'infinitely splitting, refining equivalence relations'. Its language is $L = \{E_n : n \in \omega\}$. It asserts that each E_n is an equivalence relation, E_0 consists of a single class, E_{n+1} refines E_n and in fact, partitions every E_n-class into infinitely many E_{n+1}-classes. REF(inf) is one of the standard examples of a stable, unsuperstable theory. Then:

- REF(inf) is Borel complete, in fact, it is λ-Borel complete for all infinite λ (see Definition 3.15);

- Therefore, REF(inf) is not short;
- Therefore, $\cong$ is not Borel.
- REF(inf) is grounded, i.e., $\mathrm{CSS}(\mathrm{REF}(\mathrm{inf}))_{\mathrm{sat}} = \mathrm{CSS}(\mathrm{REF}(\mathrm{inf}))_{\mathrm{ptl}}$.

REF(bin) is the theory of 'binary splitting, refining equivalence relations'. The language is also $L = \{E_n : n \in \omega\}$. The axioms of REF(bin) assert that each E_n is an equivalence relation, E_0 is trivial, each E_{n+1} refines E_n, and each E_{n+1} partitions each E_n-class into exactly two E_{n+1}-classes. REF(bin) is superstable (in fact, weakly minimal) but is not ω-stable. Then:

- T_2 (i.e., 'countable sets of reals') is Borel reducible into Mod(REF(bin)), but
- REF(bin) is short, with $||\mathrm{REF(bin)}|| = \beth_2$.
- Therefore, T_3 (i.e., 'countable sets of countable sets of reals') is not Borel reducible into Mod(REF(bin));
- Therefore, REF(bin) is not Borel complete.
- $\cong$ is not Borel.
- REF(bin) is grounded.

K is the Koerwien theory, originating in [11] and defined in Section 6. Koerwien proved that K is complete, ω-stable, eni-NDOP, and of eni-depth 2. Then:

- T_2 (i.e., 'countable sets of reals') is Borel reducible into Mod(K), and
- K is short, with $||\mathrm{K}|| = \beth_2$.
- Therefore, T_3 (i.e., 'countable sets of countable sets of reals') is not Borel reducible into Mod(K);
- Therefore, Mod(K) is not Borel complete;
- Nonetheless, $\cong$ is not Borel (this was proved by Koerwien in [11]).
- Whether K is grounded or not remains open.

TK is a 'tweaked version of K' and is also defined in Section 6. TK is also complete, ω-stable, eni-NDOP, of eni-depth 2, and is very much like the theory K, however the automorphism groups of models of TK induce a more complicated group of elementary permutations of $\mathrm{acl}(\emptyset)$ than do the automorphism groups of models of K. Then:

- TK is Borel complete, hence not short, hence $\cong$ is is not Borel.
- TK is not grounded; in fact $\mathrm{CSS}(\mathrm{TK})_{\mathrm{ptl}}$ is a proper class, while $|\mathrm{CSS}(\mathrm{TK})_{\mathrm{sat}}| = \beth_2$.
- Since $|\mathrm{CSS}(\mathrm{TK})_{\mathrm{sat}}| = \beth_2$, TK is not λ-Borel complete for sufficiently large λ.

Many of the ideas for this arose from the first author, who read and abstracted ideas about provably $\mathbf{\Delta}^{\mathbf{1}}_{\mathbf{2}}$-formulas in [3] and ideas that were discussed in Chapter 9 of [6]. Only recently, the authors became aware of work on 'pinned equivalence relations' as surveyed in e.g., [15]. The whole of this paper was written independently of the development there. In terms of notation, HC always denotes $H(\aleph_1)$, the set of all sets whose transitive closure is countable. For a forcing extension $\underline{\mathrm{V}}[G]$ of $\underline{\mathrm{V}}$, $\mathrm{HC}^{\underline{\mathrm{V}}[G]}$ denotes those sets that are hereditarily countable in $\underline{\mathrm{V}}[G]$. Throughout, we only consider set forcings, so when we write 'for $\underline{\mathrm{V}}[G]$ a forcing extension of $\underline{\mathrm{V}}\ldots$' we are quantifying over all (set) forcing notions $\underline{\mathrm{P}} \in \underline{\mathrm{V}}$ and all $\underline{\mathrm{P}}$-generic filters G.

2 A notion of cardinality for some classes of HC

Fix $\underline{\mathrm{V}}$, the universe of all sets. We develop a notion of cardinality on certain well-behaved subclasses of HC in terms of the existence or non-existence of certain injective maps. Behind the scenes, we rely heavily on the fact that all sets A in $\underline{\mathrm{V}}$ are 'potentially' elements of HC, the set of hereditarily countable sets. Specifically, if κ is the cardinality of the transitive closure of A and we take $\underline{\mathrm{P}}$ be the Levy collapsing poset $\mathrm{Coll}(\kappa^+, \omega_1)$ that collapses κ^+ to ω_1, then for any choice G of a generic filter, $A \in \mathrm{HC}^{\underline{\mathrm{V}}[G]}$.

2.1 HC-forcing invariant formulas

We begin with our principal definitions.

Definition 2.1. Suppose $\varphi(x)$ is any formula of set theory, possibly with a hidden parameter from HC.

- $\varphi(\mathrm{HC}) = \{a \in \mathrm{HC} : (\mathrm{HC}, \in) \models \varphi(a)\}$.
- If $\underline{\mathrm{V}}[G]$ is a forcing extension of $\underline{\mathrm{V}}$, then $\varphi(\mathrm{HC})^{\underline{\mathrm{V}}[G]} = \{a \in \mathrm{HC}^{\underline{\mathrm{V}}[G]} : \underline{\mathrm{V}}[G] \models \text{`}a \in \varphi(\mathrm{HC})\text{'}\}$.
- $\varphi(x)$ is *HC-forcing invariant* if, for every twice-iterated forcing extension $\underline{\mathrm{V}}[G][G']$,

$$\varphi(\mathrm{HC})^{\underline{\mathrm{V}}[G][G']} \quad \cap \quad \mathrm{HC}^{\underline{\mathrm{V}}[G]} \quad = \quad \varphi(\mathrm{HC})^{\underline{\mathrm{V}}[G]}$$

The reader is cautioned that when computing $\varphi(\mathrm{HC})^{\underline{\mathrm{V}}[G]}$, the quantifiers of φ range over $\mathrm{HC}^{\underline{\mathrm{V}}[G]}$ as opposed to the whole of $\underline{\mathrm{V}}[G]$. As examples, it follows immediately from Levy absoluteness that every Σ_1-formula, possibly with a parameter from HC, is HC-forcing invariant. As well, the class of HC-forcing invariant formulas is closed under boolean combinations.

However, for more complicated formulas, whether or not $\varphi(x)$ is HC-forcing invariant or not may well depend on the choice of set-theoretic universe. For example, assume $\underline{\mathrm{V}} = \underline{\mathrm{L}}$, and let $\varphi(x, a)$ be any non-HC-forcing invariant formula. Then the formula $\psi(x) := \varphi(x) \vee (\underline{\mathrm{V}} \neq \underline{\mathrm{L}})$ is (trivially) HC-forcing invariant whenever $\mathrm{HC} \not\subseteq \underline{\mathrm{L}}$, but it is equivalent to $\varphi(x)$ when $\mathrm{HC} \subseteq \underline{\mathrm{L}}$. Because of this, whereas we will concentrate on HC-invariant formulas in what follows, we will not say anything about the location of the formulas in any syntactic hierarchy.

Before continuing, we state three set-theoretic lemmas that form the lynchpin of our development. All three are well known. Something akin to Lemma 2.3 appears in Chapter 9 of [6], and Lemma 2.4 is reminiscent of arguments of Solovay in [14]. The key tool for all three is the Product Forcing Lemma, see e.g., Lemma 15.9 of [8], which states that given any $\underline{\mathrm{P}}_1 \times \underline{\mathrm{P}}_2$-generic filter G, if G_ℓ is the projection of G onto $\underline{\mathrm{P}}_\ell$, then $G = G_1 \times G_2$, each G_ℓ is $\underline{\mathrm{P}}_\ell$-generic, and $\underline{\mathrm{V}}[G] = \underline{\mathrm{V}}[G_1][G_2] = \underline{\mathrm{V}}[G_2][G_1]$ (i.e., G_ℓ meets every dense subset of $\underline{\mathrm{P}}_\ell$ in $\underline{\mathrm{V}}[G_{3-\ell}]$).

Lemma 2.2. *Suppose $\underline{P}_1, \underline{P}_2 \in \underline{V}$ are notions of forcing and $\varphi(x)$ is HC-forcing invariant (possibly with a hidden parameter from HC). If $A \in V$ and, for $\ell = 1, 2$, H_ℓ is $\underline{P}_\ell$-generic and $V[H_\ell] \models A \in HC$, then $\underline{V}[H_1] \models A \in \varphi(HC)$ if and only if $\underline{V}[H_2] \models A \in \varphi(HC)$. (The filters H_1 and H_2 are not assumed to be mutually generic.)*

Proof. Assume this were not the case. By symmetry, choose $p_1 \in H_1$ such that $p_1 \Vdash \check{A} \in \varphi(\mathrm{HC})$, and choose $p_2 \in H_2$ such that $p_2 \Vdash \check{A} \in \mathrm{HC} \wedge \check{A} \notin \varphi(\mathrm{HC})$. Let G be a $\underline{\mathrm{P}}_1 \times \underline{\mathrm{P}}_2$-generic filter with $(p_1, p_2) \in G$. Write $G = G_1 \times G_2$, hence $\underline{\mathrm{V}}[G] = \underline{\mathrm{V}}[G_1][G_2] = \underline{\mathrm{V}}[G_2][G_1]$. As $p_1 \in G_1$ and $p_2 \in G_2$, we have $\underline{\mathrm{V}}[G_1] \models A \in \mathrm{HC} \wedge A \in \varphi(\mathrm{HC})$ and $\underline{\mathrm{V}}[G_2] \models A \notin \varphi(\mathrm{HC})$. But, coupled with the HC-forcing invariance of φ, we obtain contradictory opinions about whether or not $\underline{\mathrm{V}}[G] \models A \in \varphi(\mathrm{HC})$. □

Lemma 2.3. *Suppose $\underline{P}_1$ and $\underline{P}_2$ are both notions of forcing in $\underline{V}$. If G is a $\underline{P}_1 \times \underline{P}_2$-generic filter and $G = G_1 \times G_2$, then $\underline{V} = \underline{V}[G_1] \cap \underline{V}[G_2]$.*

Proof. Let i_1, i_2 denote the canonical injections of $\underline{\mathrm{P}}_1, \underline{\mathrm{P}}_2$ into $\underline{\mathrm{P}}_1 \times \underline{\mathrm{P}}_2$. For a $\underline{\mathrm{P}}_1$-name τ_1, let $(\star)_{\tau_1}$ denote the following statement:

$(\star)_{\tau_1}$: For every $\underline{\mathrm{P}}_2$-name τ_2 and every $(p_1, p_2) \in \underline{\mathrm{P}}_1 \times \underline{\mathrm{P}}_2$, if $(p_1, p_2) \Vdash i_1(\tau_1) = i_2(\tau_2)$, then there is some $a \in \underline{\mathrm{V}}$ such that $p_1 \Vdash \check{a} \in \tau_1$.

It is sufficient to prove $(\star)_{\tau_1}$ for every $\underline{\mathrm{P}}_1$-name τ_1. We accomplish this by induction on the rank of the $\underline{\mathrm{P}}_1$-names. Fix an ordinal α and assume that $(\star)_{\sigma_1}$ holds for all $\underline{\mathrm{P}}_1$-names σ_1 of rank less than α. As well, choose a $\underline{\mathrm{P}}_1$-name τ_1 of rank α. To establish $(\star)_{\tau_1}$, choose τ_2 and (p_1, p_2) such that $(p_1, p_2) \Vdash i_1(\tau_1) = i_2(\tau_2)$.

Claim 1: For every $a \in \underline{\mathrm{V}}$, p_1 decides $\check{a} \in \tau_1$, i.e., either $p_1 \Vdash \check{a} \in \tau_1$ or $p_1 \Vdash \check{a} \notin \tau_1$.
Proof: Suppose not; choose $r_0, r_1 \leq p_1$ such that $r_0 \Vdash \check{a} \in \tau_1$ and $r_1 \Vdash \check{a} \notin \tau_1$.

Consider the forcing notion $\mathbb{P} := \mathbb{P}_1 \times \mathbb{P}_1 \times \mathbb{P}_2$, with canonical injections $i_0, i_1 : \mathbb{P}_1 \to \mathbb{P}$ and $i_2 : \mathbb{P}_2 \to \mathbb{P}$. Let G be $\mathbb{P}$-generic, such that $(r_0, r_1, p_2) \in G$. Then we obtain a contradiction from the following:

- $(r_0, r_1, p_2) \Vdash \check{a} \in i_0(\tau_1)$;
- $(r_0, r_1, p_2) \Vdash \check{a} \notin i_1(\tau_1)$;
- $(r_0, r_1, p_2) \Vdash i_0(\tau_1) = i_2(\tau_2) = i_1(\tau_1)$.

□

We now aim to show that $p_1 \Vdash \tau_1 \subseteq \check{\mathbb{V}}$. We begin with a set C of 'candidates'. Formally, C is the set of all 4-tuples $(r_1, r_2, \sigma_1, \sigma_2)$ that satisfy $(r_1, r_2) \in \underline{\mathrm{P}}_1 \times \underline{\mathrm{P}}_2$; for $\ell = 1, 2$ there is $q_\ell \in \underline{\mathrm{P}}_\ell$ such that $(\sigma_\ell, q_\ell) \in \tau_\ell$; and $(r_1, r_2) \Vdash i_1(\sigma_1) = i_2(\sigma_2)$. Visibly, $C \in \underline{\mathrm{V}}$. Let

$$A = \{a \in \underline{\mathrm{V}} : \text{ for some } (r_1, r_2, \sigma_1, \sigma_2) \in C,\ r_1 \Vdash \sigma_1 = \check{a}\}$$

A is a set by Replacement.

Claim 2: $p_1 \Vdash \tau_1 \subseteq \check{A}$.
Proof: Let G_1 be $\mathbb{P}_1$-generic over $\mathbb{V}$ with $p_1 \in G_1$. It suffices to show that $\mathrm{val}(\tau_1, G_1) \subseteq A$.

Let G_2 be $\mathbb{P}_2$-generic over $\mathbb{V}[G_1]$ with $p_2 \in G_2$. Let $G = G_1 \times G_2$, so by the product forcing lemma G is $\mathbb{P}_1 \times \mathbb{P}_2$-generic over $\mathbb{V}$. Since $(p_1, p_2) \in G$, we have $\mathrm{val}(i_1(\tau_1), G) = \mathrm{val}(i_2(\tau_2), G) = S$ say.

So it suffices to show that $S \subseteq A$; let $x \in S$ be arbitrary. We can write $x = \mathrm{val}(i_1(\sigma_1), G) = \mathrm{val}(i_2(\sigma_2), G)$, where $(\sigma_i, q_i) \in \tau_i$. Choose $(r_1, r_2) \in G$ such that $(r_1, r_2) \Vdash i_1(\sigma_1) = i_2(\sigma_2)$.

Then $(r_1, r_2, \sigma_1, \sigma_2) \in C$; moreover, by the inductive hypothesis, there is some $a \in \mathbb{V}$ such that $r_1 \Vdash \sigma_1 = \check{a}$. Then $a \in A$ and $x = \mathrm{val}(i_1(\sigma_1), G) = a$. □

To complete the proof of $(\star)_{\tau_1}$ let $B = \{a \in A : p_1 \Vdash \check{a} \in \tau_1\}$. It follows immediately from Claims 1 and 2 that $p_1 \Vdash \tau_1 = \check{B}$. □

Lemma 2.4. *Let $\theta(x)$ be any formula of set theory, possibly with hidden parameters from $\underline{V}$, and let $\underline{V}[G]$ be any forcing extension of $\underline{V}$. Suppose that there is some $b \in \underline{V}[G]$ such that for every forcing extension $\underline{V}[G][H]$ of $\underline{V}[G]$,*

$$\underline{V}[G][H] \models \theta(b) \wedge \exists^{=1} x \theta(x)$$

Then $b \in \underline{V}$.

Proof. Fix $\theta(x), \underline{\mathrm{V}}[G]$ and b as above. Let $\underline{\mathrm{P}} \in \underline{\mathrm{V}}$ be the forcing notion for which G is $\underline{\mathrm{P}}$-generic. Let τ be a $\mathbb{P}$-name such that $b = \mathrm{val}(\tau, G)$. Choose $p \in G$ such that

$$p \Vdash \text{“for all forcing notions } \underline{\mathrm{Q}}, \Vdash_{\underline{\mathrm{Q}}} \theta(\tau) \wedge \exists^{-1} x \theta(x)\text{”}$$

Let H be $\mathbb{P}$-generic over $\mathbb{V}[G]$ with $p \in H$. So $G \times H$ is $\mathbb{P} \times \mathbb{P}$-generic over $\mathbb{V}$. Let $i_1, i_2 : \mathbb{P} \to \mathbb{P} \times \mathbb{P}$ be the canonical injections. Then, since $(p, p) \in G \times H$, we have that

$$\mathbb{V}[G][H] \models \theta(\mathrm{val}(i_1(\tau), G \times H)) \wedge \theta(\mathrm{val}(i_2(\tau)), G \times H) \wedge \exists^{=1} x \theta(x).$$

Hence $\mathbb{V}[G][H] \models \mathrm{val}(i_1(\tau), G \times H) = \mathrm{val}(i_2(\tau), G \times H)$ and so by Lemma 2.3, $\mathrm{val}(i_1(\tau), G \times H) \in \mathbb{V}$. But $\mathrm{val}(i_1(\tau), G \times H) = b$ so we are done. □

Lemma 2.2 lends credence to the following definition.

Definition 2.5. Suppose that $\varphi(x)$ is HC-forcing invariant. Then φ_{ptl} is the class of all sets A such that $A \in \underline{\mathrm{V}}$ and, for some (equivalently, for every) forcing extension $\underline{\mathrm{V}}[G]$ of $\underline{\mathrm{V}}$, if $A \in (\mathrm{HC})^{\underline{\mathrm{V}}[G]}$, then $\underline{\mathrm{V}}[G] \models \varphi(A)$.

As motivation for the notation used in the definition above, φ_{ptl} describes the class of all $A \in \underline{\mathrm{V}}$ that are *potentially* in $\varphi(\mathrm{HC})$. We are specifically interested in those HC-forcing invariant φ for which φ_{ptl} is a **set** as opposed to a proper class.

Definition 2.6. An HC-invariant formula $\varphi(x)$ is *short* if φ_{ptl} is a set.

We begin with some easy observations. As notation, if C is a subclass of $\underline{\mathrm{V}}$, then define $\mathcal{P}(C)$ to be all sets A in $\underline{\mathrm{V}}$ such that every element of A is in C. (This definition is only novel when C is a proper class.) Similarly, $P_{\aleph_1}(C)$ is the class of all sets $A \in \mathcal{P}(C)$ that are countable (in $\underline{\mathrm{V}}$!). Let $\delta(x)$ be the formula

$$\delta(x) := \exists h[h : x \to \omega \text{ is 1-1}]$$

Given a formula $\varphi(x)$, let $\mathcal{P}(\varphi)(y)$ denote the formula $\forall x(x \in y \to \varphi(x))$ and let $\mathcal{P}_{\aleph_1}(\varphi)$ denote $\mathcal{P}(\varphi)(y) \wedge \delta(y)$.

Lemma 2.7. *1. The ptl-operator commutes with boolean combinations, i.e., if φ and ψ are both HC-forcing invariant, then $(\varphi \wedge \psi)_{ptl} = \varphi_{ptl} \cap \psi_{ptl}$ and $(\neg\varphi)_{ptl} = \underline{V} \setminus \varphi_{ptl}$.*

2. The formula $\delta(x)$ is HC-forcing invariant and $\delta_{ptl} = \underline{V}$.

3. If φ is HC-invariant, then so are both $\mathcal{P}(\varphi)$ and $\mathcal{P}_{\aleph_1}(\varphi)$. Moreover, $\mathcal{P}(\varphi_{ptl}) = (\mathcal{P}(\varphi))_{ptl} = (\mathcal{P}_{\aleph_1}(\varphi))_{ptl}$.

4. Suppose $s : \omega \to HC$ is any map such that for each n, $\varphi(x, s(n))$ is HC-forcing invariant. (Recall that HC-forcing invariant formulas are permitted to have a parameter from HC.) Then $\psi(x) := \exists n(n \in \omega \wedge \varphi(x, s(n)))$ is HC-forcing invariant and $\psi_{ptl} = \bigcup_{n\in\omega} \varphi(x, s(n))_{ptl}$.

Proof. The verification of (1) is immediate, simply by unpacking definitions.

As for (2), first note that if $h \in \underline{\mathrm{V}}$, $A \subseteq \mathrm{HC}$, and if $h : A \to \omega$ is 1-1, then $h \in \mathrm{HC}$. This observation clearly extends to any forcing extension $\underline{\mathrm{V}}[G]$, where we replace HC by $\mathrm{HC}^{\underline{\mathrm{V}}[G]}$ and conclude that $h \in \mathrm{HC}^{\underline{\mathrm{V}}[G]}$. Now, take any set $A \in \underline{\mathrm{V}}$ and choose any forcing extension $\underline{\mathrm{V}}[G]$ with $A \in \mathrm{HC}^{\underline{\mathrm{V}}[G]}$. Thus, in $\underline{\mathrm{V}}[G]$, A is countable and every element of A is contained in $\mathrm{HC}^{\underline{\mathrm{V}}[G]}$. Thus, there is an injective function $h \in \underline{\mathrm{V}}[G]$, $h : A \to \omega$. By the comments above, this $h \in \mathrm{HC}^{\underline{\mathrm{V}}[G]}$, so $\mathrm{HC}^{\underline{\mathrm{V}}[G]} \models \delta(A)$.

(3) The second equality is immediate from (1) and (2). For the first equality, begin by choosing any $A \in \mathcal{P}(\varphi)_{\text{ptl}}$. We must show that every element $a \in A$ is in φ_{ptl}. Fix any element $a \in A$. Clearly, as $A \in \underline{\mathrm{V}}$, $a \in \underline{\mathrm{V}}$. Choose any forcing extension $\underline{\mathrm{V}}[G]$ of $\underline{\mathrm{V}}$ with $A \in \mathrm{HC}^{\underline{\mathrm{V}}[G]}$. Thus

$$\mathrm{HC}^{\underline{\mathrm{V}}[G]} \models \forall x(x \in A \to \varphi(x))$$

Since $A \in \mathrm{HC}^{\underline{\mathrm{V}}[G]}$, $a \in \mathrm{HC}^{\underline{\mathrm{V}}[G]}$ as well. Thus, $\mathrm{HC}^{\underline{\mathrm{V}}[G]} \models \varphi(a)$, so $a \in \varphi_{\text{ptl}}$.

Conversely, suppose $A \in \mathcal{P}(\varphi_{\text{ptl}})$. This means that $A \in \underline{\mathrm{V}}$ and every element of A is in φ_{ptl}. Choose any forcing extension $\underline{\mathrm{V}}[G]$ of $\underline{\mathrm{V}}$ such that $A \in \mathrm{HC}^{\underline{\mathrm{V}}[G]}$. As $\mathrm{HC}^{\underline{\mathrm{V}}[G]}$ is transitive, every element $a \in A$ is also an element of $\mathrm{HC}^{\underline{\mathrm{V}}[G]}$. Thus, $\mathrm{HC}^{\underline{\mathrm{V}}[G]} \models \varphi(a)$ for every $a \in A$. That is, $\mathrm{HC}^{\underline{\mathrm{V}}[G]} \models \forall x(x \in A \to \varphi(x))$, so $A \in (\mathcal{P}(\varphi))_{\text{ptl}}$.

(4) Choose $A \in \mathrm{HC}$. It is immediate from the definition of ψ that $\mathrm{HC} \models \psi(A)$ if and only if $\mathrm{HC} \models \varphi(A, s_n)$ for some $n \in \omega$. As this equivalence relativizes to any $A \in \mathrm{HC}^{\underline{\mathrm{V}}[G]}$, both statements follow. □

2.2 Strongly definable families

In the previous subsection, we described a restricted vocabulary of formulas. Here, we discuss parameterized families of classes of $\mathrm{HC}^{\underline{\mathrm{V}}[G]}$ that are describable in this vocabulary.

Definition 2.8. Let φ be HC-forcing invariant. A family $X = (X)^{\underline{\mathrm{V}}[G]}$, indexed by the collection $\underline{\mathrm{V}}[G]$ of forcing extensions of $\underline{\mathrm{V}}$, is *strongly definable via* φ if $X^{\underline{\mathrm{V}}[G]} = \varphi(\mathrm{HC})^{\underline{\mathrm{V}}[G]}$ always. We say the family X is a *strongly definable family*, or just *strongly definable*, if it is strongly definable via some HC-forcing invariant formula φ.

We say that two HC-forcing invariant formulas φ and ψ are *persistently equivalent* if $\varphi(\mathrm{HC})^{\underline{\mathrm{V}}[G]} = \psi(\mathrm{HC})^{\underline{\mathrm{V}}[G]}$ for every forcing extension $\underline{\mathrm{V}}[G]$. Persistently equivalent formulas give rise to the same strongly definable family, and if X is strongly definable via both φ and ψ, then φ and ψ are persistently equivalent.

Definition 2.9. If X is strongly definable via φ, define X_{ptl} to be the class φ_{ptl}, i.e., all sets $A \in \underline{\mathrm{V}}$ such that $A \in (X)^{\underline{\mathrm{V}}[G]}$ for some forcing extension $\underline{\mathrm{V}}[G]$ of $\underline{\mathrm{V}}$. We call X *short* if φ is, i.e., if X_{ptl} is a set as opposed to a proper class.

Note that the condition "X is strongly definable via φ and ψ" says not only that $\varphi_{\mathrm{ptl}} = \psi_{\mathrm{ptl}}$, but that this occurs in all forcing extensions as well. Consequently it will never matter *which* formula we use as a strong definition for X, and we may work with a strongly definable X without specific reference to any particular formula φ.

We can define operations on the collection of strongly definable families. Of particular interest is the countable power set. That is, in the notation preceding Example 2.7, given any HC-forcing invariant $\varphi(x)$, the two HC-invariant formulas $\mathcal{P}(\varphi)$ and $\mathcal{P}_{\aleph_1}(\varphi)$ are persistently equal. Thus, if X is strongly definable via φ, then $\mathcal{P}(\varphi)$ and $\mathcal{P}_{\aleph_1}(\varphi)$ give rise to the same family, which we denote by $\mathcal{P}_{\aleph_1}(X)$.

We begin by enumerating several examples and easy observations that help establish our notation.

Example 2.10.

1. ω is strongly definable via the HC-forcing invariant formula $n(x) :=$"x is a natural number." Here, $\omega_{\mathrm{ptl}} = \omega$. In particular, ω is short.
2. ω_1 is strongly definable via the HC-forcing invariant formula "x is an ordinal." Here, $(\omega_1)_{\mathrm{ptl}} = ON$, the class of all ordinals. Thus, ω_1 is not short.
3. The set of reals, $\underline{\mathrm{R}} = \mathcal{P}_{\aleph_1}(\omega)$, is strongly definable via $r(x) := \mathcal{P}(n)$ in the notation before Lemma 2.7. By Lemma 2.7(3) $\underline{\mathrm{R}}_{\mathrm{ptl}} = \mathcal{P}(\omega) = \underline{\mathrm{R}}$, hence $\underline{\mathrm{R}}$ is short.
4. $\mathcal{P}_{\aleph_1}(\underline{\mathrm{R}})$, the set of countable sets of reals, is strongly definable either via $c(x) := \mathcal{P}_{\aleph_1}(r)$ or by $\mathcal{P}(r)$ (in the notation before Lemma 2.7). Via either definition, by Lemma 2.7(3), $(\mathcal{P}_{\aleph_1}(\underline{\mathrm{R}}))_{\mathrm{ptl}} = \mathcal{P}(\mathcal{P}(\omega))$, hence is short.
5. More generally, if X is short, then it follows from Lemma 2.7(3) that $\mathcal{P}_{\aleph_1}(X)_{\mathrm{ptl}} = \mathcal{P}(X_{\mathrm{ptl}})$ and so $\mathcal{P}_{\aleph_1}(X)$ is short.
6. For any $\alpha < \omega_1$, let HC_α denote the sets in HC whose rank is less than α. Then, an easy inductive argument using Lemma 2.7(3) and (4) establishes that each HC_α is strongly definable and $(\mathrm{HC}_\alpha)_{\mathrm{ptl}} = \underline{\mathrm{V}}_\alpha$. Thus, each HC_α is short.

Notation 2.11. Suppose that $X_1, \ldots, X_n$ are each strongly definable families. We say $\psi(X_1, \ldots, X_n)$ *holds persistently* if, for every forcing extension $\underline{\mathrm{V}}[G]$, we have

$$\underline{\mathrm{V}}[G] \models \psi((X_1)^{\underline{\mathrm{V}}[G]}, \ldots, (X_n)^{\underline{\mathrm{V}}[G]}).$$

We list three examples of this usage in the definition below.

Definition 2.12. Suppose that f, X and Y are each strongly definable families.

- The notation $f : X \to Y$ *persistently* means that $(f)^{\underline{V}[G]} : (X)^{\underline{V}[G]} \to (Y)^{\underline{V}[G]}$ for all forcing extensions $\underline{V}[G]$ of $\underline{V}$.
- The notation $f : X \to Y$ *is persistently injective* means that $f : X \to Y$ persistently and additionally, for all forcing extensions $\underline{V}[G]$ of $\underline{V}$, $(f)^{\underline{V}[G]} : (X)^{\underline{V}[G]} \to (Y)^{\underline{V}[G]}$ is 1-1.
- The notion $f : X \to Y$ *is persistently bijective* means $f^{-1} : X \to Y$ is strongly definable and both $f : X \to Y$ and $f^{-1} : Y \to X$ are persistently injective.

The reader is cautioned that when $f : X \to Y$ persistently (or is persistently injective), the image of f need not be strongly definable. Indeed the "image" of a strongly definable function is not well-behaved in many respects, including the lack of a surjectivity statement in the following proposition.

Proposition 2.13. *Suppose that f, X, and Y are each strongly definable.*

1. *Suppose $f : X \to Y$ persistently. Then $f_{ptl} : X_{ptl} \to Y_{ptl}$, i.e., f_{ptl} is a class function with domain X_{ptl} and image contained in Y_{ptl}.*
2. *If $f : X \to Y$ is persistently injective, then $f_{ptl} : X_{ptl} \to Y_{ptl}$ is injective as well.*
3. *If $f : X \to Y$ is persistently bijective, then $f_{ptl} : X_{ptl} \to Y_{ptl}$ is bijective.*

Proof. (1) It is obvious that f_{ptl} is a class of ordered pairs and is single-valued. As well, if $(a, b) \in f_{\mathrm{ptl}}$, then for any forcing extension $\underline{V}[G]$ of $\underline{V}$ with $(a, b) \in \mathrm{HC}^{\underline{V}[G]}$ we have $(a, b) \in (f)^{\underline{V}[G]}$ by HC-forcing invariance. Thus $\mathrm{dom}(f_{\mathrm{ptl}}) \subseteq X_{\mathrm{ptl}}$ and $\mathrm{im}(f_{\mathrm{ptl}}) \subseteq Y_{\mathrm{ptl}}$. To see that $\mathrm{dom}(f_{\mathrm{ptl}})$ is equal to X_{ptl} requires Lemma 2.4. Choose any $a \in X_{\mathrm{ptl}}$ and look at the formula $\theta(a, y)$, where $\theta(x, y)$ defines f. Let $\underline{V}[G]$ be any forcing extension of $\underline{V}$ in which $a \in (X)^{\underline{V}[G]}$. Let $b := (f)^{\underline{V}[G]}(a)$. The definition of persistence tells us that the hypotheses of Lemma 2.4 apply, hence $b \in \underline{V}$. Thus, $(a, b) \in f_{\mathrm{ptl}}$ and $a \in \mathrm{dom}(f_{\mathrm{ptl}})$.

(2) Choose $a, b, c \in \underline{V}$ such that $(a, c), (b, c) \in f_{\mathrm{ptl}}$. Choose a forcing extension $\underline{V}[G]$ of $\underline{V}$ such that $a, b, c \in \mathrm{HC}^{\underline{V}[G]}$. Thus, $(a, c), (b, c) \in (f)^{\underline{V}[G]}$. As f is persistently injective, it follows that $a = b$ holds in $\underline{V}[G]$ and hence in $\underline{V}$. So f_{ptl} is injective.

(3) From (2) we get that both f_{ptl} and $(f^{-1})_{\mathrm{ptl}}$ are (class) injections. It follows that f_{ptl} is a (class) bijection. □

We close this subsection with a characterization of surjectivity. Its proof is simply an unpacking of the definitions. However, $f : X \to Y$ being persistently surjective need not imply that the induced map $f_{\mathrm{ptl}} : X_{\mathrm{ptl}} \to Y_{\mathrm{ptl}}$ is surjective.

Lemma 2.14. *Suppose that f, X, and Y are each strongly definable via the HC-forcing invariant formulas $\theta(x, y)$, $\varphi(x)$, and $\gamma(y)$, respectively and that $f : X \to Y$ persistently. Then $f : X \to Y$ is persistently surjective if and only if the formula $\rho(y) := \exists x \theta(x, y)$ is HC-forcing invariant and persistently equivalent to $\gamma(y)$.*

2.3 A New Notion of Cardinality

Definition 2.15. Suppose X and Y are strongly definable. We say that X *is HC-reducible to* Y, written $X \leq_{\mathrm{HC}} Y$, if there is a strongly definable f such that $f : X \to Y$ is persistently injective. As notation, we write $X <_{\mathrm{HC}} Y$ if $X \leq_{\mathrm{HC}} Y$ but $Y \not\leq_{\mathrm{HC}} X$. We also write $X \sim_{\mathrm{HC}} Y$ if $X \leq_{\mathrm{HC}} Y$ and $Y \leq_{\mathrm{HC}} X$; this is apparently weaker than X and Y being in persistent bijection.

The following notion will be very useful for our applications, as it can often be computed directly. With this and Proposition 2.17 we can prove otherwise difficult non-embeddability results for $\leq_{HC}$.

Definition 2.16. Suppose X is strongly definable. The *potential cardinality* of X, denoted $\|X\|$, refers to $|X_{\mathrm{ptl}}|$ if X is short, or ∞ otherwise. By convention we say $\kappa < \infty$ for any cardinal κ.

Proposition 2.17. *Suppose X and Y are both strongly definable.*

1. *If Y is short and $X \leq_{HC} Y$, then X is short.*
2. *If $X \leq_{HC} Y$, then $\|X\| \leq \|Y\|$.*
3. *If X is short, then $X <_{HC} \mathcal{P}_{\aleph_1}(X)$.*

Proof. (1) follows immediately from (2).

(2) Choose a strongly definable f such that $f : X \to Y$ is persistently injective. Then by Proposition 2.13(2), $f_{\text{ptl}} : X_{\text{ptl}} \to Y_{\text{ptl}}$ is an injective class function. Thus, $|X_{\text{ptl}}| \leq |Y_{\text{ptl}}|$.

(3) Note that for any strongly definable X, $X \leq_{\text{HC}} \mathcal{P}_{\aleph_1}(X)$ is witnessed by the strongly definable map $x \mapsto \{x\}$. For the other direction, suppose by way of contradiction that X is short, but $\mathcal{P}_{\aleph_1}(X) \leq_{\text{HC}} X$. Also, $(\mathcal{P}_{\aleph_1}(X))_{\text{ptl}} = \mathcal{P}(X_{\text{ptl}})$ by Lemma 2.7(3). Thus, by (2), we would have that $|\mathcal{P}(X_{\text{ptl}})| \leq |X_{\text{ptl}}|$, which contradicts Cantor's theorem since X_{ptl} is a set. □

Using the fact that $(\text{HC}_\beta)_{\text{ptl}} = \underline{\text{V}}_\beta$, the following Corollary is immediate.

Corollary 2.18. *If X is strongly definable and $\|X\| \leq \beth_\alpha$ for some $\alpha < \omega_1$, then $HC_{\omega+\alpha+1} \not\leq_{HC} X$.*

2.4 Quotients

We begin with the obvious definition.

Definition 2.19. A pair (X, E) is a *strongly definable quotient* if both X and E are strongly definable and persistently, E is an equivalence relation on X.

There is an immediate way to define a reduction of two quotients:

Definition 2.20. Let (X, E) and (Y, F) be strongly definable quotients. Say $(X, E) \leq_{\text{HC}} (Y, F)$ if there is a strongly definable f such that all of the following hold persistently:

- f is a subclass of $X \times Y$.
- The E-saturation of $\text{dom}(f)$ is X. That is, for every $x \in X$, there is an $x' \in X$ and $y' \in Y$ where xEx' holds and $(x', y') \in f$.
- f induces a well-defined injection on equivalence classes. That is, if (x, y) and (x', y') are in f, then xEx' holds if and only if yFy' does.

Define $(X, E) <_{\text{HC}} (Y, F)$ and $(X, E) \sim_{\text{HC}} (Y, F)$ in the natural way.

Unfortunately, because $X_{\text{ptl}}/E_{\text{ptl}}$ need not be $(X/E)_{\text{ptl}}$, the "obvious" way to define $(X/E)_{\text{ptl}}$ doesn't work in general. For our purposes, we can restrict to a more well-behaved subclass.

Definition 2.21. A *representation* of a strongly definable quotient (X, E) is a pair f, Z of strongly definable families such that $f : X \to Z$ is persistently surjective and persistently,

$$\forall a, b \in X[E(a, b) \Leftrightarrow a = b]$$

We say that (X, E) is *representable* if it has a representation.

In the case that (X, E) is representable, the class of E-classes is strongly definable in a sense – we equate it with the representation. For this reason we will also say Z is a representation of (X, E). Note that if $f_1 : X \to Z_1$ and $f_2 : X \to Z_2$ are two representations of (X, E), then there is a persistently bijective, strongly definable $h : Z_1 \to Z_2$. This observation implies the following definition is well-defined.

Definition 2.22. If (X, E) is a representable strongly definable quotient, then define $||(X, E)|| = ||Z||$ for some (equivalently, for all) Z such that there is a representation $f : (X, E) \to Z$.

The following lemma can be proved by a routine composition of maps.

Lemma 2.23. *Suppose (X, E) and (Y, E') are strongly definable quotients, with representations $f : (X, E) \to Z$ and $g : (Y, E') \to Z'$. Suppose h is a witness to $(X, E) \leq_{HC} (Y, E')$. Then the induced function $h^* : Z \to Z'$ is strongly definable, persistently injective, and witnesses $Z \leq_{HC} Z'$.*

We close this section with an observation about restrictions of representations.

Lemma 2.24. *Suppose $f : (X, E) \to Z$ is a representation and $Y \subseteq X$ is strongly definable and persistently E-saturated. Let E' and $g = f\restriction_Y$ be the restrictions of E and f, respectively, to Y. Then the image $g(Y)$ is strongly definable, and so $g : (Y, E') \to g(Y)$ is a representation.*

Proof. We show that $\varphi(z) := \exists y(y \in Y \wedge g(y) = z)$ is HC-forcing invariant. Fix any $z \in HC$ and assume $z \in \varphi(HC)^{\underline{\mathrm{V}}[G]}$ for some forcing extension $\underline{\mathrm{V}}[G]$ of $\underline{\mathrm{V}}$. Choose a witness $y \in (Y)^{\underline{\mathrm{V}}[G]}$ such that $g(y) = z$ in $\underline{\mathrm{V}}[G]$. As $f(y) = z$, we conclude $z \in (Z)^{\underline{\mathrm{V}}[G]}$ and hence $z \in Z^{\underline{\mathrm{V}}}$ as Z is strongly definable. So, choose $y^* \in (X)^{\underline{\mathrm{V}}}$ with $f(y^*) = z$. Thus, in $\underline{\mathrm{V}}[G]$, $E(y, y^*)$ holds. As Y is persistently E-saturated, $y^* \in (Y)^{\underline{\mathrm{V}}[G]}$. Since Y is strongly definable, we conclude $y^* \in (Y)^{\underline{\mathrm{V}}}$, so y^* witnesses that $z \in (g(Y))^{\underline{\mathrm{V}}}$. As this argument relativizes to any forcing extension, we conclude that $\varphi(z)$ is HC-forcing invariant. □

All of the examples we work with will be representable, where the representations are Scott sentences. Therefore this simple definition of $\|(X, E)\|$ will suffice completely for our purposes. In the absence of a representation, one can still define $\|(X, E)\|$ using the notion of *pins*; see for instance [15] for a thorough discussion.

3 Connecting Potential Cardinality with Borel Reducibility

The standard framework for Borel reducibility of invariant classes is the following. Fix two countable languages L and L'. Let X_L and $X_{L'}$ denote the sets of L and L'-structures with universe ω. If we endow each set with the usual logic topology, then each of X_L and $X_{L'}$ becomes a (standard) Polish space. Moreover, if Φ is a sentence of $L_{\omega_1,\omega}$ then $\mathrm{Mod}(\Phi)$ is a Borel subset of X_L and dually for $L'_{\omega_1,\omega}$ sentences Φ'. Classically, a Borel reduction from $(\mathrm{Mod}(\Phi), \cong) \to (\mathrm{Mod}(\Phi), \cong)$ is a Borel map $f : \mathrm{Mod}(\Phi) \to \mathrm{Mod}(\Phi')$ such that, for all $M, N \in \mathrm{Mod}(\Phi)$, $M \cong N$ if and only if $f(M) \cong f(N)$.

There are two noteworthy observations about this setup. First, the ambient spaces are standard Polish spaces, and the second is that a Borel reduction is really a mapping on the quotient spaces $\mathrm{Mod}(\Phi)/\cong$ and $\mathrm{Mod}(\Psi)/\cong$. As each of the languages is countable, we can view each of the spaces X_L and $X_{L'}$ of models with universe ω as classes in HC. It is then easily checked that for any sentence Φ, $\mathrm{Mod}(\Phi)$ is a strongly definable class. In order to invoke the machinery developed in the previous section, we seek a representation of the quotient $(\mathrm{Mod}(\Phi), \cong)$. Happily, we will see below that the classical notion of a canonical Scott sentence, when coupled with Karp's completeness theorem for sentences of $L_{\omega_1,\omega}$, yields a representation.

3.1 Canonical Scott sentences

For what follows, we need the notion of a canonical Scott sentence of any infinite L-structure, regardless of cardinality. The definition below is in both Barwise [1] and Marker [13].

Definition 3.1. Suppose L is countable and M is any infinite L-structure, say of power κ. For each $\alpha < \kappa^+$, define an $L_{\kappa^+,\omega}$ formula $\varphi^{\overline{a}}_\alpha(\overline{x})$ for each finite $\overline{a} \in M^{<\omega}$ as follows:

- $\varphi^{\overline{a}}_0(\overline{x}) := \bigwedge\{\theta(\overline{x}) : \theta \text{ atomic or negated atomic and } M \models \theta(\overline{a})\}$;
- $\varphi^{\overline{a}}_{\alpha+1}(\overline{x}) := \varphi^{\overline{a}}_\alpha(\overline{x}) \ \wedge \ \bigwedge\left\{\exists y\, \varphi^{\overline{a},b}_\alpha(\overline{x}, y) : b \in M\right\} \ \wedge \forall y \bigvee\left\{\varphi^{\overline{a},b}_\alpha(\overline{x}, y) : b \in M\right\}$;
- For α a non-zero limit, $\varphi^{\overline{a}}_\alpha(\overline{x}) := \bigwedge\left\{\varphi^{\overline{a}}_\beta(\overline{x}) : \beta < \alpha\right\}$.

Next, let $\alpha^*(M) < \kappa^+$ be least such that for all finite $\overline{a}, \overline{a}'$ from M,

$$\forall \overline{x}[\varphi^{\overline{a}}_{\alpha^*(M)}(\overline{x}) \leftrightarrow \varphi^{\overline{a}'}_{\alpha^*(M)}(\overline{x})] \Longrightarrow \forall \overline{x}[\varphi^{\overline{a}}_{\alpha^*(M)+1}(\overline{x}) \leftrightarrow \varphi^{\overline{a}'}_{\alpha^*(M)+1}(\overline{x})]$$

Finally, put $\mathrm{css}(M) := \varphi^{\emptyset}_{\alpha^*(M)} \wedge \bigwedge \left\{ \forall \overline{x}[\varphi^{\overline{a}}_{\alpha^*(M)}(\overline{x}) \to \varphi^{\overline{a}}_{\alpha^*(M)+1}(\overline{x})] : \overline{a} \in M^{<\omega} \right\}$.

For what follows, it is crucial that the choice of $\mathrm{css}(M)$ really is canonical. In particular, in the infinitary clauses forming the definition of $\mathrm{tp}^{\overline{a}}_{\alpha+1}(\overline{x})$, we consider the conjunctions and disjunctions be taken over *sets* of formulas, as opposed to sequences. In particular, we ignore the multiplicity of a formula inside the set. By our conventions about working wholly in ZFC, countable languages and sentences of $L_{\infty,\omega}$ are sets, and in particular $\mathrm{css}(M)$ is a set. Moreover, if M is countable, then $\mathrm{css}(M) \in HC$.

We summarize the well-known, classical facts about canonical Scott sentences with the following:

Fact 3.2. *Fix a countable language L.*

1. *For every L-structure M, $M \models \mathit{css}(M)$; and for all L-structures N, $M \equiv_{\infty,\omega} N$ if and only if $\mathit{css}(M) = \mathit{css}(N)$ if and only if $N \models \mathit{css}(M)$.*
2. *If M is countable, then $\mathit{css}(M) \in HC$; and*
3. *css is highly absolute, i.e.,*
   - *If $M \in HC$, then $(\mathit{css}(M))^{HC} = \mathit{css}(M)$; and*
   - *If $M \in \underset{\sim}{V}$ and $\underset{\sim}{V}[G]$ is any forcing extension of $\underset{\sim}{V}$, then $(\mathit{css}(M))^{\underset{\sim}{V}[G]} = \mathit{css}(M)$.*
4. *If M and N are both countable, then $M \cong N$ if and only if $\mathit{css}(M) = \mathit{css}(N)$ if and only if $N \models \mathit{css}(M)$.*

Our primary interest in canonical Scott sentences is that they give rise to representations of classes of L-structures. The key to the representability is Karp's Completeness theorem for sentences of $L_{\omega_1,\omega}$. We begin by considering $\mathrm{CSS}(L)$, the set of all canonical Scott sentences of structures in X_L, the set of L-structures with universe ω.

Lemma 3.3. *Fix a countable language L. Then:*

1. *$\mathit{CSS}(L)$ is strongly definable via the formula $\varphi(y) := \exists M(M \in X_L \wedge \mathit{css}(M) = y)$;*
2. *The strongly definable function $\mathit{css} : X_L \to \mathit{CSS}(L)$ is persistently surjective;*
3. *$\mathit{css} : X_L \to \mathit{CSS}(L)$ is a representation of the quotient $(X_L, \cong)$.*

Proof. (1) We need to verify that $\varphi(y)$ is HC-forcing invariant. Suppose $\sigma \in \mathrm{HC}$ and $\underset{\sim}{V}[G]$ is a forcing extension of $\underset{\sim}{V}$. If $\sigma \in \varphi(\mathrm{HC})^{\underset{\sim}{V}[G]}$, then there is some $M \in X_L^{\underset{\sim}{V}[G]}$ such that $\mathrm{css}(M) = \sigma$. Thus, in $\underset{\sim}{V}[G]$, σ is consistent because it is satisfiable. But, as consistency is absolute, coupled with $\sigma \in \mathrm{HC}$, yields that in $\underset{\sim}{V}$, σ is a consistent sentence of $L_{\omega_1,\omega}$. So, by Karp's completeness theorem, there is some (countable) $N \in (X_L)^{\underset{\sim}{V}}$ such that $N \models \sigma$. In $\underset{\sim}{V}[G]$, $N \models \sigma$, which is $\mathrm{css}(M)$. So $\mathrm{css}(N) = \sigma$ in $\underset{\sim}{V}[G]$, and hence in $\underset{\sim}{V}$. As this argument readily relativizes to any forcing extension, φ is HC-forcing invariant.

(2) follows from (1) and Lemma 2.14.

(3) is immediate by (2) and Fact 3.2(2). □

In most of our applications, we are interested in strongly definable subclasses of X_L that are closed under isomorphism. For any sentence Φ of $L_{\omega_1,\omega}$, because $\mathrm{Mod}(\Phi)$ is a Borel subset of X_L, $\mathrm{Mod}(\Phi)$ is also strongly definable, as is the restriction map (also denoted by css) $\mathrm{css} : \mathrm{Mod}(\Phi) \to \mathrm{HC}$.

Definition 3.4. For Φ any sentence of $L_{\omega_1,\omega}$, $\mathrm{CSS}(\Phi) = \{\mathrm{css}(M) : M \in \mathrm{Mod}(\Phi)\} \subseteq HC$.

Proposition 3.5. *Fix any sentence $\Phi \in L_{\omega_1,\omega}$ in a countable vocabulary. Then $\mathit{css} : \mathit{Mod}(\Phi) \to \mathit{CSS}(\Phi)$ is a representation of the quotient $(\mathit{Mod}(\Phi), \cong)$.*

Proof. As $\mathrm{Mod}(\Phi)$ is strongly definable, this follows immediately from Lemmas 3.3 and 2.24. □

Definition 3.6. Let Φ be any sentence of $L_{\omega_1,\omega}$ in a countable vocabulary. We say that Φ is *short* if $\mathrm{CSS}(\Phi)_{\mathrm{ptl}}$ is a set (as opposed to a proper class). If Φ is short, let the *potential cardinality* of Φ, denoted $||\Phi||$, be the (usual) cardinality of $\mathrm{CSS}(\Phi)_{\mathrm{ptl}}$; otherwise let it be ∞.

It follows from Proposition 3.5 and Definitions 2.15 and 2.22 that

$$||\Phi|| = ||(\mathrm{Mod}(\Phi), \cong)|| = ||\mathrm{CSS}(\Phi)|| = |\mathrm{CSS}(\Phi)_{\mathrm{ptl}}|.$$

In order to understand the class $\mathrm{CSS}(\Phi)_{\mathrm{ptl}}$, note that if $\varphi \in \mathrm{CSS}(\Phi)_{\mathrm{ptl}}$, then $\varphi \in \underline{\mathrm{V}}$ and is a sentence of $L_{\infty,\omega}$. If we choose any forcing extension $\underline{\mathrm{V}}[G]$ of $\underline{\mathrm{V}}$ for which $\varphi \in \mathrm{HC}^{\underline{\mathrm{V}}[G]}$, then $\underline{\mathrm{V}}[G] \models$ '$\varphi \in L_{\omega_1,\omega}$' and there is some $M \in \mathrm{HC}^{\underline{\mathrm{V}}[G]}$ such that $V[G] \models$ '$M \in \mathrm{Mod}(\Phi)$ and $\mathrm{css}(M) = \varphi$'. Thus, we refer to elements of $\mathrm{CSS}(\Phi)_{\mathrm{ptl}}$ as being *potential canonical Scott sentences* of a model of Φ. In particular, every element of $\mathrm{CSS}(\Phi)_{\mathrm{ptl}}$ is *potentially satisfiable* in the sense that it is satisfiable in some forcing extension $\underline{\mathrm{V}}[G]$ of $\underline{\mathrm{V}}$. There is a proof system for sentences of $L_{\infty,\omega}$ for which a sentence is consistent if and only if it is potentially satisfiable as defined above. Thus, when we say 'φ implies ψ', we mean that for any forcing extension $\underline{\mathrm{V}}[G]$, every model of φ is a model of ψ.

A more basic question to ask is to determine the image of the class function css when restricted to the class of models of Φ. As notation, let $\mathrm{CSS}(\Phi)_{\mathrm{sat}}$ denote the class of $\{\mathrm{css}(M) : M \in \underline{\mathrm{V}} \text{ and } M \models \Phi\}$. This choice of notation is clarified by the following easy lemma.

Lemma 3.7. *$CSS(\Phi)_{sat} \subseteq CSS(\Phi)_{ptl}$.*

Proof. Choose any $\varphi \in \mathrm{CSS}(\Phi)$ and choose any $M \in \underline{\mathrm{V}}$ such that $M \models \Phi$ and $\mathrm{css}(M) = \varphi$. Then clearly, $\varphi \in \underline{\mathrm{V}}$. As well, choose a forcing extension $\underline{\mathrm{V}}[G]$ in which M is countable. Then, in $\underline{\mathrm{V}}[G]$, there is some $M' \in \mathrm{Mod}(\Phi)$ (i.e., where the universe of M' is ω) such that $M' \cong M$. As the class function css is absolute, $(\mathrm{css}(M'))^{\underline{\mathrm{V}}[G]} = \varphi$ by Fact 3.2. Thus, $\varphi \in \mathrm{CSS}(\Phi)_{\mathrm{ptl}}$. □

To summarize, elements of $\mathrm{CSS}(\Phi)_{\mathrm{ptl}}$ are called *potential* canonical Scott sentences, whereas elements of $\mathrm{CSS}(\Phi)_{\mathrm{sat}}$ are *satisfiable*. This suggests a property of the sentence Φ.

Definition 3.8. A sentence $\Phi \in L_{\omega_1,\omega}$ (or a complete first-order theory T) is *grounded* if $\mathrm{CSS}(\Phi)_{\mathrm{sat}} = \mathrm{CSS}(\Phi)_{\mathrm{ptl}}$, i.e., if every potential canonical Scott sentence is satisfiable.

As a trivial example, if T is $\aleph_0$-categorical, then as all models of T are back-and-forth equivalent, $\mathrm{CSS}(T)_{\mathrm{ptl}}$ is a singleton, hence T is grounded. In Section 5 we show that both of the theories REF(bin) and REF(inf) are grounded, but in Section 6 we prove that the theory TK is not grounded.

Next, we show that a Borel reduction between invariant classes yields a strongly definable map between the associated canonical Scott sentences.

Fact 3.9. *Suppose Φ and Φ' are sentences of $L_{\omega_1,\omega}$ and $L'_{\omega_1,\omega}$ respectively. If there is a Borel reduction $f : (Mod(\Phi), \cong) \to (Mod(\Phi'), \cong)$ then there is a strongly definable $f^* : CSS(\Phi) \to CSS(\Phi')$ between canonical Scott sentences that is persistently injective. Hence $CSS(\Phi) \leq_{HC} CSS(\Phi')$.*

Proof. It is a standard theorem, see e.g., [9], that the graph of f is Borel. So f is naturally a strongly definable family of Borel sets. By Lemma 2.23, it suffices to show that f witnesses $(\mathrm{Mod}(\Phi), \cong) \leq_{\mathrm{HC}} (\mathrm{Mod}(\Phi'), \cong)$, which amounts to showing that f remains well-defined and injective on isomorphism classes in every forcing extension. But this is a $\mathbf{\Pi^1_2}$ statement in codes for f, Φ, Φ', and thus is absolute to forcing extensions by Shoenfield's Absoluteness Theorem. □

The following Theorem is simply a distillation of our previous results.

Theorem 3.10. *Let Φ and Ψ be sentences of $L_{\omega_1,\omega}$, possibly in different countable vocabularies.*

1. *If Ψ is short, while Φ is not short, then $(Mod(\Phi), \cong)$ is not Borel reducible to $(Mod(\Psi), \cong)$.*

2. *If $||\Psi|| < ||\Phi||$, then $(Mod(\Phi), \cong)$ is not Borel reducible to $(Mod(\Psi), \cong)$.*

Proof. (1) follows immediately from (2).

(2) Suppose $f : (\mathrm{Mod}(\Phi), \cong) \to (\mathrm{Mod}(\Psi), \cong)$ were a Borel reduction. Then by Fact 3.9 we would obtain a strongly definable $f^* : \mathrm{CSS}(\Phi) \to \mathrm{CSS}(\Psi)$ that is persistently injective, meaning $||\mathrm{CSS}(\Phi)|| \leq ||\mathrm{CSS}(\Psi)||$. This contradicts Proposition 2.17(2). □

3.2 Consequences of $\cong$ being Borel

Although our primary interest is classes $\mathrm{Mod}(\Phi)$ where $\cong$ is not Borel, in this brief subsection we see the consequences of $\cong$ being Borel.

Theorem 3.11. *The following are equivalent for a sentence $\Phi \in L_{\omega_1,\omega}$ in countable vocabulary.*

1. *The relation of $\cong$ on $Mod(\Phi)$ is a Borel subset of $Mod(\Phi) \times Mod(\Phi)$;*

2. *For some $\alpha < \omega_1$, $CSS(\Phi) \subseteq HC_\alpha$;*

3. *For some $\alpha < \omega_1$, $CSS(\Phi)$ is persistently contained in HC_α;*

4. *$CSS(\Phi)_{ptl}$ is contained in $\underline{V}_\alpha$ for some $\alpha < \omega_1$.*

Proof. To see the equivalence of (1) and (2), first note that in both conditions we are only considering models of Φ with universe ω and the canonical Scott sentence of such objects. In particular, neither condition involves passing to a forcing extension. However, it is a classical result that $\cong$ is Borel if and only if there is a uniform bound on the heights of canonical Scott sentences, which is equivalent to stating that there is a bound in the HC_α hierarchy.

For (2) implies (3), note that the formula $\exists M : M \models \Phi \wedge \mathrm{css}(M) \notin \mathrm{HC}_\alpha$ is a Σ_1 formula in the parameters Φ, α and so is absolute to forcing extensions.

That (3) implies (4) and (4) implies (2) follow directly from Example 2.10(6). □

We obtain an immediate corollary to this. Let $I_{\infty,\omega}(\Phi)$ denote the cardinality of a maximal set of pairwise $\equiv_{\infty,\omega}$-inequivalent models $M \in \underline{\mathrm{V}}$ (of any cardinality) with $M \models \Phi$. If no maximal set exists, we write $I_{\infty,\omega}(\Phi) = \infty$. By Fact 3.2 and Lemma 3.7, $I_{\infty,\omega}(\Phi) = |\mathrm{CSS}(\Phi)_{\mathrm{sat}}| \leq ||\Phi||$.

Corollary 3.12. *Let Φ be any sentence in $L_{\omega_1,\omega}$ in a countable vocabulary such that $\cong$ is a Borel subset of $Mod(\Phi) \times Mod(\Phi)$. Then*

1. *Φ is short; and*

2. *$I_{\infty,\omega}(\Phi) < \beth_{\omega_1}$.*

Proof. Assume that $\cong$ is a Borel subset of $\mathrm{Mod}(\Phi) \times \mathrm{Mod}(\Phi)$. By Theorem 3.11(3), $\mathrm{CSS}(\Phi)_{\mathrm{ptl}} \subseteq \underline{\mathrm{V}}_\alpha$ for some $\alpha < \omega_1$ and hence is a set. Thus, Φ is short. So $I_{\infty,\omega}(\Phi) = |\mathrm{CSS}(\Phi)_{\mathrm{sat}}| \leq |\mathrm{CSS}(\Phi)_{\mathrm{ptl}}| \leq |\underline{\mathrm{V}}_\alpha| < \beth_{\omega_1}$ □

We remark that the implication in Corollary 3.12 does not reverse. In Sections 5 and 6 we show that both of the complete theories REF(bin) and K are short, but on countable models of either theory, $\cong$ is not Borel.

3.3 Maximal Complexity

In this subsection, we recall two definitions of maximality. The first, Borel completeness, is from Friedman-Stanley [3].

Definition 3.13. Let $\Phi \in L_{\omega_1,\omega}$. The quotient $(\mathrm{Mod}(\Phi), \cong)$ is *Borel complete* if every $(\mathrm{Mod}(\Psi), \cong)$ is Borel reducible to $(\mathrm{Mod}(\Phi), \cong)$.

Corollary 3.14. *If $(Mod(\Phi), \cong)$ is Borel complete, then Φ is not short.*

Proof. Let $L = \{\leq\}$ and let Ψ assert that $\leq$ is a linear ordering. As $(\mathrm{Mod}(\Phi), \cong)$ is Borel complete, there is a Borel reduction $f : (\mathrm{Mod}(\Psi), \cong) \to (\mathrm{Mod}(\Phi), \cong)$. However, it is easily proved that for distinct ordinals $\alpha \neq \beta$, the L-structures $(\alpha, \leq)$ and $(\beta, \leq)$ are $\equiv_{\infty\omega}$-inequivalent models of Ψ, hence have distinct canonical Scott sentences. Thus, $\mathrm{CSS}(\Psi)_{\mathrm{sat}}$ is a proper class, and hence so is $\mathrm{CSS}(\Psi)_{\mathrm{ptl}}$ by Lemma 3.7. So Φ cannot be short by Theorem 3.10(1). □

If one is only interested in classes of countable models, then the Borel complete classes are clearly maximal with respect to Borel reducibility. As any invariant class of countable structures has a natural extension to a class of uncountable structures, one can ask for more. The following definitions from [12] generalize Borel completeness to larger cardinals λ. To see that it is a generalization, recall that among countable structures, isomorphism is equivalent to back-and-forth equivalence, and that for structures of size λ, $\equiv_{\lambda^+,\omega}$-equivalence is also equivalent to back-and-forth equivalence. Consequently, 'Borel complete' in the sense of Definition 3.13 is equivalent to '$\aleph_0$-Borel complete' in Definition 3.15. So, 'Φ is λ-Borel complete for all infinite λ' implies Φ Borel complete. However, in Section 6 we will see that the theory TK is Borel complete, but is not λ-Borel complete for large λ.

Definition 3.15.

- For $\lambda \geq \aleph_0$, let $\mathrm{Mod}_\lambda(\Phi)$ denote the class of models of Φ with universe λ.
- Toplogize $\mathrm{Mod}_\lambda(\Phi)$ by declaring that $\mathcal{B} := \{U_{\theta(\overline{\alpha})} : \theta(\overline{x}) \text{ is quantifier free and } \overline{\alpha} \in \lambda^{<\omega}\}$ is a sub-basis, where $U_{\theta(\overline{\alpha})} = \{M \in \mathrm{Mod}_\lambda(\Phi) : M \models \theta(\overline{\alpha})\}$.
- A function $f : \mathrm{Mod}_\lambda(\Phi) \to \mathrm{Mod}_\lambda(\Psi)$ is a *λ-Borel embedding* if
  - the inverse image of every (sub)-basic open set is λ-Borel; and
  - For $M, N \in \mathrm{Mod}_\lambda(\Phi)$, $M \equiv_{\infty,\omega} N$ if and only if $f(M) \equiv_{\infty,\omega} f(N)$.
- $(\mathrm{Mod}_\lambda(\Phi), \equiv_{\infty,\omega})$ is *λ-Borel reducible* to $(\mathrm{Mod}_\lambda(\Psi), \equiv_{\infty,\omega})$ if there exists a λ-Borel embedding $f : \mathrm{Mod}_\lambda(\Phi) \to \mathrm{Mod}_\lambda(\Psi)$; and
- Φ is *λ-Borel complete* if every $(\mathrm{Mod}_\lambda(\Psi), \equiv_{\infty,\omega})$ is λ-Borel reducible to $(\mathrm{Mod}_\lambda(\Phi), \equiv_{\infty,\omega})$.

As examples, the class of graphs (directed or undirected) is λ-Borel complete for all infinite λ. In [12] it is proved that the class of subtrees of $\lambda^{<\omega}$ is λ-Borel complete, and more recently the second author has proved that the class of linear orders is λ-Borel complete for all λ.

3.4 Jumps and products

In this subsection we recall two procedures – the *jump* and the *product* – and use them to define a sequence $\langle T_\alpha : \alpha \in \omega_1 \rangle$ of complete, first order theories for which the potential cardinality is strictly increasing.

Definition 3.16. Suppose L is a countable relational language and $\Phi \in L_{\omega_1,\omega}$. The *jump of* Φ, written $J(\Phi)$, uses the language $L' = L \cup \{E\}$, where E is a binary relation not appearing in L. $J(\Phi)$ states that E is an equivalence relation with infinitely many classes, each of which is a model of Φ. If $R \in L$ and $\emptyset x$ is a tuple not all from the same E-class, then $R(\emptyset x)$ is defined to be false, so that the models are independent.

The notion of a jump was investigated in [3], where it was shown that if E is a Borel equivalence relation with more than one class, then $E <_B J(E)$. We give a partial generalization of this in Proposition 3.17(3) – if $\Phi \in L_{\omega_1,\omega}$ is short, then $\|\Phi\| < \|J(\Phi)\|$, so $\Phi <_B J(\Phi)$. Using the theory of pins [15], one can use essentially the same proof to give a true generalization: if (X, E) is short with more than one E-class (without regard to representability), then $\|(X, E)\| <_{\mathrm{HC}} \|(X^\omega, J(E))\|$, so in particular $E <_B J(E)$.

The following Proposition lists the basic properties of the jump operation.

Proposition 3.17. *Let Φ and Ψ be $L_{\omega_1,\omega}$-sentences in countable relational languages.*

1. *If Φ is a complete first order theory, so is $J(\Phi)$.*
2. *If Φ is grounded, so is $J(\Phi)$.*
3. *If Φ is short, then $J(\Phi)$ is also short. In particular, if $\|\Phi\|$ is infinite, $\|J(\Phi)\| = 2^{\|\Phi\|}$. If $2 \leq \|\Phi\| < \aleph_0$, $\|J(\Phi)\| = \aleph_0$. If $\|\Phi\| = 1$, then $\|J(\Phi)\| = 1$.*
4. *The jump is monotone: if $\Phi \leq_B \Psi$, then $J(\Phi) \leq_B J(\Psi)$.*
5. *It is always true that $\Phi \leq_B J(\Phi)$; if Φ is short and not $\aleph_0$-categorical, then $\Phi <_B J(\Phi)$.*

Note that here and throughout, we use $\Phi \leq_B \Psi$ as a shorthand for $(Mod(\Phi), \cong) \leq_B (Mod(\Psi), \cong)$, and similarly with $<_B$ and $\sim_B$.

Proof. (1) That the jump is first-order is clear. Completeness follows from a standard Ehrenfeucht-Fraïsse argument.

(2) Let $\Psi \in \mathrm{CSS}(J(\Phi))_{\mathrm{ptl}}$, and let $\underline{\mathrm{V}}[G]$ be some forcing extension in which Ψ is hereditarily countable. Let $M \models \Psi$ be the unique countable model of Ψ in $\underline{\mathrm{V}}[G]$. Let X be the set of E-classes in M, and for each $x \in X$, let Ψ_x be the canonical Scott sentence of x, viewed as an L-structure. Let m_x be the number of E-classes of M which are isomorphic to x as L-structures, if this number is finite; if infinite, let $m_x = \omega$.

The set of pairs $S = \{(\Psi_x, m_x) : x \in X\}$ depends only on the isomorphism type of M, so is uniquely definable from Ψ. By Lemma 2.4, $S \in \underline{\mathrm{V}}$, although it may no longer be countable. Since Φ is grounded, each Ψ_x has a model in $\underline{\mathrm{V}}$, so let N be the model coded from S – for each pair (Ψ, m) in S, give m distinct E-classes, each of which is a model of Ψ. Since each Ψ is $\aleph_0$-categorical, $M \cong N$ in any sufficiently large forcing extension of $\underline{\mathrm{V}}[G]$, so $N \models \Psi$, as desired.

(3) We assume $\|\Phi\|$ is infinite; the finite cases are similar and trivial, respectively. First, let $X \subset \mathrm{CSS}(\Phi)_{\mathrm{ptl}}$ be arbitrary. Let $\underline{\mathrm{V}}[G]$ be a forcing extension in which X is hereditarily countable, and let M_X be a countable model of $J(\Phi)$ such that each E-class of M_X is a model of some $\Psi \in X$, and each $\Psi \in X$ is represented exactly once. Since each Ψ is $\aleph_0$-categorical, M_X is determined up to isomorphism by these constraints. Therefore, $\Psi_X = \mathrm{CSS}(M_X)$ is determined entirely by X, so $\Psi_X \in \underline{\mathrm{V}}$ by Lemma 2.4. Furthermore, if $X \neq Y$, then $M_X \ncong M_Y$, so $\Psi_X \neq \Psi_Y$, so $\|J(\Phi)\| \geq 2^{\|\Phi\|}$.

For the other direction, let $\Psi \in \mathrm{CSS}(J(\Phi))$, let $\underline{\mathrm{V}}[G]$ be a forcing extension in which Ψ is hereditarily countable, and let $M_\Psi \models \Psi$ be the unique countable model. Let $\mathcal{C}$ be the set of E-classes of M_Ψ, and for each $C \in \mathcal{C}$, let Ψ_C be $\mathrm{CSS}(C)$, where we consider the E-class C as an L-structure and a model of Ψ. For each C, let $m(C)$ be the number of equivalence classes C' where $\Psi_{C'} = \Psi_C$; as usual, if this number is infinite, let $m(C) = \omega$. Since M_Ψ is determined up to isomorphism by Ψ, the set $S_\Psi = \{(\Psi_C, m(C))$ is determined entirely by Ψ. Therefore, $S_\Psi \in \underline{\mathrm{V}}$ by Lemma 2.4. If $\Psi \neq \Psi'$, $M_\Psi \ncong M_{\Psi'}$, so $S_\Psi \neq S_{\Psi'}$. Therefore, this defines an injection from $\|J(\Phi)\|$ into $(\omega + 1)^{\|\Phi\|}$. Since $\|\Phi\|$ is infinite, this last is equal to $2^{\|\Phi\|}$, completing the proof.

(4) is trivial; (5) follows from (3) and the fact that if $\|\Phi\| > \|\Psi\|$ then $\Phi \not\leq_B \Psi$. □

Another important operation is the *product*:

Definition 3.18. Suppose I is a countable set and for each i, Φ_i is a sentence of $L_{\omega_1,\omega}$ in the countable relational language L_i. The *product of the* Φ_i, denoted $\prod_i \Phi_i$, is in the language $L = \{U_i : i \in I\} \cup \bigcup_i L_i$, where we consider the L_i to be disjoint and the U_i to be unary predicates not appearing in any L_i.

Then $\prod_i \Phi_i$ states that the U_i are disjoint, that the elements of U_i form a model of Φ_i when viewed as an L_i-structure, and that if $R \in L_i$ and $\varnothing x$ is not all from U_i, then $R(\varnothing x)$ is false, so that the models are independent. If I is finite, we also require that each element is in some U_i.

The proofs of the corresponding facts for products are quite similar to those for the jump, so we omit the proofs:

Proposition 3.19. *Let $\{\Phi_i : i \in I\}$ and $\{\Psi_j : j \in J\}$ be countable sets of $L_{\omega_1,\omega}$-sentences in countable relational languages.*

1. *If each Φ_i is complete and first order, so is $\prod_i \Phi_i$.*
2. *If each Φ_i is grounded, so is $\prod_i \Phi_i$.*
3. *If each Φ_i is short, then $\prod_i \Phi_i$ is short. In particular, $\| \prod_i \Phi_i \| = \kappa \cdot \prod_i \|\Phi_i\|$, where κ is either one, if I is finite, or $\aleph_0$, if I is infinite.*
4. *The product is monotone: if $f : I \to J$ is an injection and $\Phi_i \leq_B \Psi_{f(i)}$ for all $i \in I$, then $\prod_i \Phi_i \leq_B \prod_j \Psi_j$.*
5. *It is always true that for all $i \in I$, $\Phi_i \leq_B \prod_i \Phi_i$. If additionally, for all $i \in I$, there is an $i' \in J$ where $\Phi_i <_B \Phi_{i'}$, then for all $i \in I$, $\Phi_i <_B \prod_i \Phi_i$.*

Note that if we use $L_{\omega_1,\omega}$ to add to our definition of product that the U_i are exhaustive, the product of first-order theories may not be first order, but the statement of (3) improves and we can lose reference to κ. This can also be achieved by working in a multisorted first-order logic, which preserves (1) and fixes (3). We do not express a preference among these choices.

We close this section by defining a set of concrete benchmarks: Note that these are essentially the same as the $\mathcal{I}_\alpha$ discussed in [3] and the $\cong_\alpha$ in [7] as well as the hierarchy of theories discussed in [10].

Definition 3.20. T_0 is the theory of $(\underline{\mathbb{Z}}, S)$, where S is the graph of the successor function.

For each countable α, $T_{\alpha+1} = J(T_\alpha)$. If α is a limit ordinal, $T_\alpha = \prod_{\beta<\alpha} T_\beta$.

The theories T_α, in one form or another, have been studied extensively in [3] and [7], where they were denoted $\mathcal{I}_\alpha$ and $\cong_\alpha$, respectively; we use the notation T_α to emphasize that these are first-order theories. We quickly summarize the properties which are relevant to us. First and foremost, they are concrete enough to work with. They are also all *Borel*; that is, $\cong_{T_\alpha}$ is Borel for every α. Finally, they are cofinal; if $\cong_\Phi$ is Borel, then $\Phi \leq_B T_\alpha$ for some α. Therefore, characterizing the relationship between some Φ and the T_α is a reasonably strong way to gauge the complexity of isomorphism for Φ.

This gives us the basic properties for our "benchmark" theories:

Corollary 3.21. *Each T_α is a complete first-order theory which is short and grounded. For finite α, $\|T_\alpha\| = \beth_\alpha$; for infinite α, $\|T_\alpha\| = \beth_{\alpha+1}$. In particular, $T_\alpha <_B T_\beta$ for all $\alpha < \beta < \omega_1$.*

Proof. We first check $\alpha = 0$. Clearly T_0 has exactly $\aleph_0$ models up to back-and-forth equivalence. If $M \models T_0$ in some $\underline{\mathbb{V}}[G]$, then either M has finite dimension n, so is isomorphic to a model in $\underline{\mathbb{V}}$, or M has infinite dimension, so is back-and-forth equivalent to the countable model $(\omega \times \underline{\mathbb{Z}}, S)$ in $\underline{\mathbb{V}}$. Either way, this simultaneously shows that T_0 is grounded and $||T_0|| = \aleph_0$, completing the proof.

The rest of the proof goes immediately by induction, using Propositions 3.17 and 3.19. Groundedness follows from part (2). Size counting (and therefore shortness) follows from part (3), which is immediate at successor stages but the limits require an argument. So let α be a limit ordinal. Then $\|T_\alpha\| = \aleph_0 \cdot \prod_\beta \|T_\beta\|$. By the inductive hypothesis and usual tricks about cardinal multiplication, this is equal to $\prod_\beta \beth_\beta = \prod_\beta 2^{\beth_\beta}$, this last being equal to $2^{\sum_\beta \beth_\beta} = 2^{\beth_\alpha} = \beth_{\alpha+1}$, completing the proof.

The strictness of the ascending chain follows from induction and part (5). □

4 Compact group actions

In this brief section we use the technology of canonical Scott sentences and representability to analyze the effect of a continuous action of a compact group on a Polish space X. In particular, we show that the quotient of $\mathcal{P}_{\aleph_1}(X)$ by the diagonal action of G is representable. We also show that if the group is abelian, we can bound the potential cardinality of the representation. In Section 6 we use these results to analyze the models of the theory K and to contrast K with TK.

Throughout this subsection, we fix a recursively presented Polish space X, a strongly definable group G, and a strongly definable $\cdot : G \times X \to X$ that is persistently a continuous action. We also fix an effective countable basis

$$\mathcal{B}_n = \{U_i^n : i \in \omega\}$$

of X^n for each $n \geq 1$.

The action of G on X naturally gives diagonal actions on both X^n and $\mathcal{P}(X)$ defined by $g \cdot \overline{a} = \langle g \cdot a : a \in \overline{a} \rangle$ and $g \cdot A = \{g \cdot a : a \in A\}$, respectively. Clearly, the diagonal action of G takes countable subsets of X to countable subsets. Let $\sim^G$ be the equivalence relation on $\mathcal{P}_{\aleph_1}(X)$ given by

$$A \sim^G B \quad \text{if and only if} \quad g \cdot A = B \text{ for some } g \in G$$

In order to understand the quotient $(\mathcal{P}_{\aleph_1}(X), \sim^G)$, we begin with one easy lemma that uses the fact that G is compact. This lemma is the motivation for the first order language we define below.

Lemma 4.1. *If $A, B \in \mathcal{P}_{\aleph_1}(X)$, then $A \sim^G B$ if and only if there is a bijection $\sigma : A \to B$ satisfying $\overline{a} \sim^G h(\overline{a})$ for all $\overline{a} \in A^{<\omega}$.*

Proof. If $g \cdot A = B$, then $\sigma := g \restriction_A$ is as desired. For the converse, fix such a σ; we will show there is $g \in G$ extending σ. Let $\{a_n : n \in \omega\}$ be an enumeration of A, and for each n, let $\overline{a}_n$ be the tuple $a_0 \dots a_{n-1}$ and let $C_n \subseteq G$ be the set of all $g \in G$ with $g \cdot \overline{a}_n = \sigma(\overline{a}_n)$. C_n is closed since the action is continuous and C_n is nonempty by hypothesis. Since G is compact, $C = \bigcap_n C_n$ is nonempty, and clearly any $g \in C$ has $g \cdot A = B$. □

We define a first-order language L and a class of L-structures that encode this information. Put $L := \{R_i^n : i \in \omega, n \geq 1\}$, where each R_i^n is an n-ary relation. Let M_X be the L-structure with universe all of X, with each R_i^n interpreted by

$$M_X \models R_i^n(\overline{a}) \quad \text{if and only if} \quad G\overline{a} \cap U_i^n = \emptyset$$

As notation, let $\mathrm{qf}_n(\overline{a})$ denote the quantifier-free type of $\overline{a} \in X^n$. It is easily seen that

$$\mathrm{qf}_n(\overline{a}) = \mathrm{qf}_n(\overline{b}) \quad \text{if and only if} \quad G\overline{a} = G\overline{b}$$

As well, note that every $g \in G$ induces an L-automorphism of M_X given by $a \mapsto g \cdot a$. These two observations imply that M_X has a certain homogeneity – For $\overline{a}, \overline{b} \in X^n$, $\mathrm{qf}_n(\overline{a}) = \mathrm{qf}_n(\overline{b})$ if and only if there is an automorphism of M_X taking $\overline{a}$ to $\overline{b}$.

For $\overline{a}, \overline{b} \in X^n$ the relation $\overline{a} \sim^G \overline{b}$ is absolute between $\underline{\mathrm{V}}$ and any forcing extension $\underline{\mathrm{V}}[H]$. In addition, as the bases of our topologies on X^n are recursively presented, it follows from Σ_1-absoluteness that if $\Gamma \in \underline{\mathrm{V}}$ is any quantifier-free n-type, then for any forcing extension $\underline{\mathrm{V}}[H]$ of $\underline{\mathrm{V}}$, Γ is realized in $(M_X)^{\underline{\mathrm{V}}[H]}$ if and only if Γ is realized in M_X.

As notation, call an L-structure $N \in HC$ *nice* if it is isomorphic to a substructure of M_X. Let $\mathcal{W}$ consist of all nice L-structures.

Lemma 4.2. *An L-structure $N \in HC$ is nice if and only if for every $n \geq 1$, every quantifier-free n-type realized in N is realized in M_X.*

Proof. Left to right is obvious. For the converse, choose any $N \in HC$ for which every quantifier-free n type realized in N is realized in M_X. We construct an L-embedding of N into M_X via a "forth" construction using the homogeneity of M_X. Enumerate the universe of $N = \{a_n : n \in \omega\}$ and let $\overline{a}_n$ denote $\langle a_i : i < n\rangle$. Assuming $f_n : \overline{a}_n \to M_X$ has been defined, choose any $\overline{b} \in X^{n+1}$ such that $\mathrm{qf}_{n+1}(\overline{a}_{n+1}) = \mathrm{qf}_{n+1}(\overline{b})$. Write $\overline{b}$ as $\overline{b}_n b^*$. As $\mathrm{qf}_n(\overline{b}_n) = \mathrm{qf}_n(f_n(\overline{a}_n))$, there is an automorphism σ of M_X with $\sigma(\overline{b}_n) = f_n(\overline{a}_n)$. Then define f_{n+1} to extend f_n and satisfy $f_{n+1}(a_n) = \sigma(b^*)$. □

Define a map $f : \mathcal{P}_{\aleph_1}(X) \to \mathcal{W}$ by $A \mapsto M_A$, the substructure of M_X with universe A.

Our first goal is the following Theorem.

Theorem 4.3. *For any recursively presented Polish space X and any strongly definable, compact group G with a strongly definable, persistently continuous action, we have that:*

1. *Both $\mathcal{W}$ and f are strongly definable;*
2. *For all $A, B \in \mathcal{P}_{\aleph_1}(X)$, $A \sim^G B$ if and only if $f(A) \cong f(B)$, persistently;*
3. *The canonical Scott sentence map $css : (\mathcal{W}, \cong) \to CSS(\mathcal{W})$ is a representation;*
4. *The quotient $(\mathcal{P}_{\aleph_1}(X), \sim^G)$ is representable via the composition map $css \circ f$ that takes $A \mapsto css(M_A)$.*

Proof. (1) It is obvious that f is strongly definable.

That $\mathcal{W}$ is strongly definable follows from Lemma 4.2 and the absoluteness results mentioned above. In particular, for any L-structure $N \in \underline{\mathrm{V}}$ that is not nice, there is some n and $\overline{a} \in N^n$ such that $\Gamma := \mathrm{qf}_n(\overline{a})$ is not realized in M_X. But then, in any forcing extension $\underline{\mathrm{V}}[H]$, $(M_X)^{\underline{\mathrm{V}}[H]}$ does not realize Γ, so N is not nice in $\underline{\mathrm{V}}[H]$. As this argument relativizes to any forcing extension, $\mathcal{W}$ is strongly definable.

For (2), if $A \sim^G B$, then any $g \in G$ that satisfies $g \cdot A = B$ induces a bijection between A and B such that $\overline{a} \sim^G g \cdot \overline{a}$ for all $\overline{a} \in A^{<\omega}$. As this implies $G\overline{a} = G(g \cdot \overline{a})$, $\mathrm{qf}_n(\overline{a})$ in M_A is equal to $\mathrm{qf}_n(g \cdot \overline{a})$ in M_B. Thus, the action by g induces an isomorphism of the L-structures M_A and M_B. Conversely, suppose $\sigma : M_A \to M_B$ is an L-isomorphism. Then $\sigma(\overline{a}) \sim^G \overline{a}$ for every $\overline{a} \in A^{<\omega}$, so $A \sim^G B$ by Lemma 4.1.

(3) As $\mathcal{W}$ is strongly definable, this follows immediately from Lemmas 3.3 and 2.24.

(4) follows immediately from (1), (2), and (3). □

As a consequence of Theorem 4.3, $||(\mathcal{P}_{\aleph_1}(X), \sim^G)||$ is defined. For an arbitrary compact group action, this quotient need not be short. Indeed, Theorem 6.8 gives an example where it is not. However, if we additionally assume that G is abelian, then we will see below that $||(\mathcal{P}_{\aleph_1}(X), \sim^G)|| \leq \beth_2$. The reason for this stark discrepancy is due to the comparative simplicity of abelian group actions. In particular, if an abelian group G acts transitively on a set S, then S is essentially an affine copy of $G/\mathrm{Stab}(a)$, where $\mathrm{Stab}(a)$ is the subgroup of G stabilizing any particular $a \in S$. The following Lemma is really a restatement of this observation.

Lemma 4.4. *Suppose that an abelian group G acts on a set X. Then for every $n \geq 1$, if three n-tuples $\overline{a}, \overline{b}, \overline{c} \in X^n$ satisfy $\overline{a} \sim^G \overline{b} \sim^G \overline{c}$ and $\overline{a}\overline{b} \sim^G \overline{a}\overline{c}$, then $\overline{b} = \overline{c}$.*

Proof. Let $g \in G$ be such that $g \o a = \o b$. Choose $h \in G$ such that $h(\o a \o b) = \o a \o c$. Then in particular, $h \o a = \o a$ and $h \o b = \o c$. From this, $gh \o a = \o b$ and $hg \o a = \o c$. But $gh = hg$, so $\o b = \o c$, as desired. □

We will show that $\|(\mathcal{W}, \cong)\| \leq \beth_2$ by showing that each Scott sentence in the representation is from $L_{\beth_1^+\omega}$, and then using the fact that there are at most $\beth_2$ such sentences. In the case that the Scott sentence is satisfiable, this means the model has size at most $\beth_1$, so this is perhaps unsurprising. We will accomplish this complexity bound by a type-counting argument; here is the notion of type we will use.

If φ is a canonical Scott sentence – that is, $\varphi \in \mathrm{CSS}(L)_{\mathrm{ptl}}$ – then let $S_n^\infty(\varphi)$ be the set of all canonical Scott sentences in the language $L' = L \cup \{c_0, \ldots, c_{n-1}\}$ which imply φ. We will refer to elements of $S_n^\infty(\varphi)$ as types – infinitary formulas with free variables $x_0, \ldots, x_{n-1}$, resulting from replacing each c_i with a new

variable x_i not otherwise appearing in the formula. It is equivalent to define $S_n^\infty(\varphi)$ by forcing – if $\underline{\mathrm{V}}[H]$ makes φ hereditarily countable and $M \in \underline{\mathrm{V}}[H]$ is the unique countable model of φ, then $S_n^\infty(\varphi)$ is the set $\{\mathrm{css}(M, \emptyset a) : \emptyset a \in M^n\}$. Evidently this set depends only on the isomorphism class of M, so by the usual argument with Lemma 2.4, this set is in $\underline{\mathrm{V}}$.

Proposition 4.5. *Suppose φ is a canonical Scott sentence in a language of size at most κ, and for all n, $|S_n^\infty(\varphi)| \leq \kappa$, where κ is an infinite cardinal. Then φ is a sentence of $L_{\kappa^+\omega}$.*

Proof. We use the precise syntactic definition of Scott formulas from Definition 3.1. For a moment, pass to a forcing extension $\underline{\mathrm{V}}[G]$ in which φ is hereditarily countable, and let M be its unique countable model. For each ordinal α, let $S_n^\alpha(\varphi)$ be the set $\{\varphi_\alpha^{\emptyset a}(\emptyset x) : \emptyset a \in M^n\}$, recalling that $S_n^\infty(\varphi)$ is the set $\{\mathrm{css}(M, \emptyset a) : \emptyset a \in M^n\}$. Since M is (persistently above $\underline{\mathrm{V}}[G]$) unique up to isomorphism, and since these sets are unchanged by passing to an isomorphic image of M, $S_n^\alpha(\varphi)$ and $S_n^\infty(\varphi)$ are in $\underline{\mathrm{V}}$ and depend only on φ. Moreover, there is a natural surjection $\pi_n^\alpha : S_n^\infty(\varphi) \to S_n^\alpha(\varphi)$ taking $\mathrm{css}(M, \emptyset a)$ to $\varphi_\alpha^{\emptyset a}(\emptyset x)$. Also by Lemma 2.4, the π_n^α are in $\underline{\mathrm{V}}$.

Let α^* be the Scott rank of M. Again, this is invariant under isomorphism, so depends only on φ. For any two sentences $\psi, \tau \in S_n^\infty(\varphi)$, let $d(\psi, \tau)$ be the least $\alpha < \alpha^*$ where $\pi_n^{\alpha+1}(\psi) \neq \pi_n^{\alpha+1}(\tau)$; if there is no such α, then $\psi = \tau$, so say $d(\psi, \tau) = \alpha^*$. It is immediate from the construction of Scott rank that if $\alpha \leq \alpha^*$, there are Scott sentences ψ and τ of some arity where $d(\psi, \tau) = \alpha$; hence $d : \bigcup_n (S_n^\infty(\varphi))^2 \to \alpha^* + 1$ is surjective. Further, d depends only on the formulas themselves, so by Lemma 2.4, $d \in \underline{\mathrm{V}}$.

The rest of the proof takes place in $\underline{\mathrm{V}}$. Because π_n^α is surjective and $|S_n^\infty(\varphi)| \leq \kappa$, $|S_n^\alpha(\varphi)| \leq \kappa$ for all α. Similarly, $|\bigcup_n (S_n^\infty(\varphi))^2| \leq \kappa$ and d is surjective, so $|\alpha^* + 1| \leq \kappa$. By induction we show that for all $\alpha \leq \alpha^* + 1$, $S_n^\alpha(\varphi) \subseteq L_{\kappa^+\omega}$. The base case is trivial, since there are only κ atomic formulas. The step follows from the fact that $|S_n^\alpha(\varphi)| \leq \kappa$, and the limit follows from the fact that $\alpha^* < \kappa^+$, so in both cases we need only take conjunctions and disjunctions of κ formulas at a time.

Observe that φ is precisely the following:

$$\pi_0^{\alpha^*}(\varphi) \wedge \bigwedge \left\{ \forall \emptyset x \left(\pi_n^{\alpha^*}(\varphi^*)(\emptyset x) \to \pi_n^{\alpha^*+1}(\varphi^*)(\emptyset x) \right) : n \in \omega, \varphi^* \in S_n^\infty(\varphi) \right\}$$

Since $S_n^\alpha(\varphi) \subseteq L_{\kappa^+\omega}$ for all α and n, and since they all have size at most κ, φ is in $L_{\kappa^+\omega}$, as desired. □

The following holds by a straightforward induction on the complexity of formulas:

Lemma 4.6. *For all infinite cardinals κ and languages L of size at most κ, there are exactly 2^κ different $L_{\kappa^+\omega}$ formulas.*

Now we can prove our theorem. Recall that $\sim_G$ is the diagonal equivalence relation on $\mathcal{P}_{\aleph_1}(X)$, induced by the diagonal action of G.

Theorem 4.7. *Let X be a strongly definable, recursively presented Polish space and let G be a strongly definable, compact abelian group acting continuously on X. Suppose all this holds persistently. Then $\|(\mathcal{P}_{\aleph_1}(X), \sim_G)\| \leq \beth_2$.*

Proof. We use Proposition 4.5 to show that $\mathrm{CSS}(\mathcal{W})_{\mathrm{ptl}} \subseteq L_{\beth_1^+\omega}$; then by Lemma 4.6, we have that $|\mathrm{CSS}(\mathcal{W})_{\mathrm{ptl}}| \leq \beth_2$, as desired. So let $\varphi \in \mathrm{CSS}(\mathcal{W})_{\mathrm{ptl}}$ be arbitrary; it is enough to show that $|S_n^\infty(\varphi)| \leq \beth_1$.

For each n, let $\mathrm{qf}_n(\varphi)$ be the set of quantifier-free n-types which are consistent with φ. We have a surjective map $\pi_n : S_n^\infty(\varphi) \to \mathrm{qf}_n(\varphi)$ sending $\psi(\emptyset x)$ to the set of quantifier-free formulas in $\emptyset x$ which it implies. For any $p \in \mathrm{qf}_n(\varphi)$, let $S_n^\infty(\varphi, p)$ be $\pi^{-1}(p)$, the set of $\psi \in S_n^\infty(\varphi)$ where $\pi_n(\psi) = p$. Since the language is countable, $|\mathrm{qf}_n(\varphi)| \leq \beth_1$. Thus it is sufficient to show that for all p, $|S_n^\infty(\varphi, p)| \leq \beth_1$.

Now we take advantage of the fact that G is abelian:

Claim: Suppose $p^* \in \mathrm{qf}_{2n}(\varphi)$ is such that $p^* \restriction_{[0,n)} = p^* \restriction_{[n,2n)} = p$. Suppose that $\psi^*, \tau^* \in S_{2n}^\infty(\varphi)$ *both* complete p^*. Further, suppose $\psi^* \restriction_{[0,n)} = \tau^* \restriction_{[0,n)}$. Then $\psi^* = \tau^*$.

Proof: Pass to a forcing extension $\underline{\mathrm{V}}[H]$ in which φ is hereditarily countable, and let M be its unique countable model. By Theorem 4.3, we may assume $M = M_A$ for some $A \in \mathcal{P}_{\aleph_1}(X)^{\underline{\mathrm{V}}[H]}$. Choose some tuples $(\emptyset a_0 \emptyset a_1)$

and $(\varnothing b_0 \varnothing b_1)$ from M^{2n} where $\mathrm{css}(M, \varnothing a_0 \varnothing a_1) = \psi^*$ and $\mathrm{css}(M, \varnothing b_0 \varnothing b_1) = \tau^*$. By assumption $\mathrm{css}(M, \varnothing a_0) = \mathrm{css}(M, \varnothing b_0)$, so we may assume $\varnothing a_0 = \varnothing b_0$. Since all of the tuples $\varnothing b_0$, $\varnothing a_1$, and $\varnothing b_1$ have the same quantifier-free type, they are in the same G-orbit, and similarly with $\varnothing b_0 \varnothing a_1$ and $\varnothing b_0 \varnothing b_1$. Thus Lemma 4.4 applies directly to the triple $(\varnothing b_0, \varnothing b_1, \varnothing a_1)$, so in particular $\varnothing b_1 = \varnothing a_1$. Thus $\psi^* = \mathrm{css}(M, \varnothing b_0 \varnothing a_1) = \mathrm{css}(M, \varnothing b_0 \varnothing b_1) = \tau^*$, as desired. □

Then for *any* $\psi \in S_n^\infty(\varphi, p)$, define $\Gamma(\psi)$ to be the set of all $p^* \in \mathrm{qf}_{2n}(\varphi)$ such that $p^* \restriction_{[0,n)} = p^* \restriction_{[n,2n)} = p$ such that for *some* $\psi^* \in S_{2n}^\infty(\varphi, p^*)$, $\psi^* \restriction_{[0,n)} = \psi$.

By the Claim, if $p^* \in \Gamma(\psi)$, there is a *unique* $\psi^* \in S_{2n}^\infty(\varphi)$ where $\pi_{2n}(\psi^*) = p^*$ and $\psi^* \restriction_{[0,n)} = \psi$. So define $F(p^*)$ to be $\psi^* \restriction_{[n,2n)}$. Evidently $|\Gamma(\psi)| \leq \beth_1$, so it is enough to show that $F : \Gamma(\psi) \to S_n^\infty(\varphi)$ is surjective.

But this is almost immediate. Fix any $\tau \in S_n^\infty(\varphi)$ and let $\underline{\mathbf{V}}[H]$ be a forcing extension in which φ is hereditarily countable, and let M be its unique countable model; as before, we may assume $M = M_A$ for some $A \in \mathcal{P}_{\aleph_1}(X)^{\underline{\mathbf{V}}[H]}$. Choose any $\varnothing a \in A^n$ where $\mathrm{css}(M, \varnothing a) = \psi$ and any $\varnothing b \in A^n$ where $\mathrm{css}(M, \varnothing a) = \tau$. Finally, let p^* be the quantifier-free type of $\varnothing a \varnothing b$ in M. Clearly $p^* \in \Gamma(\psi)$ and $F(p^*) = \tau$. □

This theorem will be crucial in Section 6.

5 Refining Equivalence Relations

We begin by defining an incomplete first-order theory REF. Its language is $L = \{E_n : n \in \omega\}$ and its axioms posit:

- Each E_n is an equivalence relation;
- E_0 has a single equivalence class; that is, we consider xE_0y to be universally true;
- For all n, E_{n+1} refines E_n; that is, every E_n-class is a union of E_{n+1}-classes.

The theory REF is very weak, which makes the generality of the following proposition surprising.

Proposition 5.1. *REF is grounded.*

Proof. We begin with an analysis of an arbitrary model M of REF. As notation, for any $a \in M$ and $n \in \omega$, let $[a]_n$ denote the equivalence class of a, i.e., $\{b \in M : M \models E_n(a, b)\}$. As the equivalence relations refine each other, the classes $T(M) = \{[a]_n : a \in M, n \in \omega\}$ form an ω-tree, ordered by $[a]_n \leq [b]_m$ if and only if $n \leq m$ and $[b]_m \subseteq [a_n]$. Next, let E_∞ be the equivalence relation given by $E_\infty(a, b)$ if and only if $E_n(a, b)$ for every $n \in \omega$. Let $[a]_\infty$ be the E_∞-class of a. Then M/E_∞ can be construed as a subset of the branches $[T(M)]$ of $T(M)$. As we are interested in determining models up to back-and-forth equivalence (as opposed to isomorphism), the following definition is natural.

For each $a \in M$, let the *color of* a, $c(a) \in (\omega + 1) \setminus \{0\}$ be given by

$$c(a) = \begin{cases} |[a]_\infty| & \text{if } [a]_\infty \text{ is finite} \\ \omega & \text{if } [a]_\infty \text{ is infinite} \end{cases}$$

Next, we describe some expansions of M to larger languages. For each $n \in \omega$, let $L_n = L \cup \{U_i : i \leq n\}$, where the U_i's are distinct unary predicates. Given any $M \models \mathrm{REF}$, $n \in \omega$, and $a \in M$, let $M_n(a)$ denote the L_n-structure $(M, [a]_0, \ldots, [a]_n)$, i.e., where each predicate U_i is interpreted as $[a]_i$.

We now exhibit some invariants, which we term the *data of* M, written $D(M)$ which we will see only depend on the $\equiv_{\infty,\omega}$-equivalence class of M.

For each $n \in \omega$, let

$$I_n(M) = \{\mathrm{css}(M_n(a)) : a \in M\}$$

. We combine the sets $I_n(M)$ into a tree $(I(M), \leq)$ where $I(M) = \bigcup_{n \in \omega} I_n(M)$ and, for $\sigma_n \in I_n(M)$ and $\psi_m \in I_m(M)$, we say $\sigma_n \leq \psi_m$ if and only if $n \leq m$ and $\psi_m \vdash \sigma_n$. That is, if the reduct of any model of

ψ_m to L_n is a model of σ_n. As back-and-forth equivalence is preserved under reducts, it follows that when $n \leq m$, every $\sigma_m \in I_m$ is above precisely one $\sigma_n \in I_n$. Thus, $(I(M), \leq)$ is an ω-tree.

Continuing, for each $n > 0$ and $\sigma_n \in I_n(M)$, let the multiplicity of σ_n, $\text{mult}_M(\sigma_n) \in (\omega + 1) \setminus \{0\}$, be given by: $\text{mult}_M(\sigma_n) = k < \omega$ if k is maximal such that there are elements $\{b_i : i < k\} \subseteq M$ such that

$$\bigwedge_{i<j<k} \Big[E_{n-1}(b_i, b_j) \wedge \neg E_n(b_i, b_j) \Big] \wedge \bigwedge_{i<k} \text{css}(M_n(b_i)) = \sigma_n$$

and let $\text{mult}_M(\sigma_n) = \omega$ if there is an infinite family $\{b_i : i < \omega\}$ as above.

Now, each $a \in M$ induces a *canonical sequence* $Seq_M(a) := \langle \text{css}(M_n(a)) : n \in \omega \rangle$, which is clearly a branch through the tree $I(M)$. Let $Seq(M) = \{Seq_M(a) : a \in M\}$. Clearly, $Seq(M) \subseteq [I(M)]$, the set of branches of $I(M)$. Finally, for any $s \in Seq(M)$, we define the *color spectrum of* s, $Sp_M(s) = \{c(a) : Seq_M(a) = s\}$. Thus, each $Sp_M(s)$ is a non-empty subset of $(\omega + 1) \setminus \{0\}$.

Define the *data of* M, $D(M) := \langle (I(M), \leq), \text{mult}_M, Seq(M), Sp_M \rangle$.

Claim 1: For any $M, N \models \text{REF}$, $M \equiv_{\infty,\omega} N$ if and only if $D(M) = D(N)$.
Proof: First, note that since $[a]_0 = M$ for every $a \in M$, $M \equiv_{\infty,\omega} N$ is trivially equivalent to $M_0(a) \equiv_{\infty,\omega} N_0(b)$ for every $a \in M$ and $b \in N$. Thus, if $D(M) = D(N)$, then as the trees $(I(M), \leq)$ and $(I(N), \leq)$ are equal, they have the same root, so $\text{css}(M_0(a)) = \text{css}(N_0(b))$ for some/every $a \in M, B \in N$. So $M \equiv_{\infty,\omega} N$.

The forward direction is a bit more tedious. For what follows, observe the following easy facts: If $(M, a) \equiv_{\infty,\omega} (N, b)$, then

- For each $n \in \omega$, $M_n(a) \equiv_{\infty,\omega} N_n(b)$, so $\text{css}(M_n(a)) = \text{css}(N_n(b))$; and
- $c(a) = c(b)$.

Now suppose $M \equiv_{\infty,\omega} N$. We first argue that for each $n \in \omega$, the sets $I_n(M)$ and $I_n(N)$ are equal and that $\text{mult}_M(\sigma_n) = \text{mult}_N(\sigma_n)$. For $n = 0$, this follows from the fact that $\text{css}(M_0(a))$ is the canonical Scott sentence of M for every choice of $a \in M$. For $n > 0$, choose any $\sigma_n \in I_n$ and choose any $a \in M$ such that $\sigma_n = \text{css}(M_n(a))$. As $M \equiv_{\infty,\omega} N$, choose $b \in N$ such that $(M, a) \equiv_{\infty,\omega} (N, b)$. By the Fact above, it follows that $\text{css}(N_n(b)) = \text{css}(M_n(a)) = \sigma_n$, so $I_n(M) \subseteq I_n(N)$ and the other direction follows by symmetry. Thus, $I_n(M) = I_n(N)$. As for multiplicity, choose any $\sigma_n \in I_n$ and suppose that $\{b_i : i < k\} \subseteq M$ witnesses that $\text{mult}_M(\sigma_n) \geq k$. Then choose $\{c_i : i < k\} \subseteq N$ such that

$$(M, b_0, \ldots, b_{k-1}) \equiv_{\infty,\omega} (N, c_0, \ldots, c_{k-1})$$

It follows that $\text{mult}_N(\sigma_n) \geq k$ as well. By symmetry, it follows that $\text{mult}_M(\sigma_n) = \text{mult}_N(\sigma_n)$ as well.

To see that $Seq(M) = Seq(N)$, choose any $s \in Seq(M)$. Choose $a \in M$ such that $s = Seq_M(a)$ and choose $b \in N$ such that $(M, a) \equiv_{\infty,\omega} (N, b)$. By the Fact above, we have that $\text{css}(N_n(b)) = \text{css}(M_n(a))$ for every $n \in \omega$, hence $Seq_N(b) = Seq_M(a) = s$. Thus, $Seq(M) \subseteq Seq(N)$, and hence $Seq(M) = Seq(N)$ by symmetry.

Finally, fix any $s \in Seq(M)$ and choose any color $k \in Sp_M(s)$. Choose $a \in M$ such that $Seq_M(a) = s$ and $c(a) = k$. Choose $b \in N$ such that $(M, a) \equiv_{\infty,\omega} (N, b)$. By the Fact, $Seq_N(b) = s$ and $c(b) = k$. So, by symmetry, $Sp_M(s) = Sp_N(s)$. □

To begin the proof of groundedness, choose any $\sigma \in \text{CSS}(\text{REF})_{\text{ptl}}$. Choose any forcing extension $\underline{\text{V}}[G]$ of $\underline{\text{V}}$ in which $\sigma \in HC^{\underline{\text{V}}[G]}$ and hence $\sigma \in \text{CSS}(\text{REF})^{\underline{\text{V}}[G]}$. Choose any model $M \in \underline{\text{V}}[G]$ with $\text{css}(M) = \sigma$. Working in $\underline{\text{V}}[G]$, compute $D(M)$, the data of M. However, in light of Claim 1, $D(M)$ only depends on σ. As σ is fixed, for the remainder of the argument we write

$$D = \langle (I, \leq), \text{mult}, Seq, Sp \rangle$$

to represent these data. We argue that although $M \in \underline{\text{V}}[G]$, $D \in \underline{\text{V}}$. To see this, choose any forcing extension $\underline{\text{V}}[G][H]$ of $\underline{\text{V}}[G]$ and choose any $M' \in \underline{\text{V}}[G][H]$ such that $M' \models \sigma$. As $\text{css}(M) = \text{css}(M')$, it follows from

Fact 3.2 that $M \equiv_{\infty,\omega} M'$. Thus, by applying Claim 1 in $\underline{\mathrm{V}}[G][H]$, $D(M) = D(M')$. So, it follows from Lemma 2.4 that $D \in \underline{\mathrm{V}}$.

To complete the proof of the Proposition, we work in $\underline{\mathrm{V}}$ and 'unpack' the data D and construct an L-structure $N \in \underline{\mathrm{V}}$ such that in $\underline{\mathrm{V}}[G]$, $M \equiv_{\infty,\omega} N$. Once we have this, as $\sigma = \mathrm{css}(M)$, it follows that $N \models \sigma$ in $\underline{\mathrm{V}}[G]$ and hence in $\underline{\mathrm{V}}$. Thus, N witnesses that $\sigma \in \mathrm{CSS}(\mathrm{REF})_{\mathrm{sat}}$. That is, the proof of groundedness will be finished once we establish the following Claim.

Claim 2: Working in $\underline{\mathrm{V}}$, there is an L-structure N such that $\underline{\mathrm{V}}[G] \models$ '$N \equiv_{\infty,\omega} M$'.
Proof: Before beginning the 'unpacking' of D, we note some connections between M and D that are not part of the data. First, there is a surjective tree homomorphism $h : T(M) \cup M/E_\infty \to I \cup Seq$ given by $h([a]_n) = \mathrm{css}(M_n(a))$ for $n \in \omega$ and $h([a]_\infty) = \langle \mathrm{css}(M_n(a)) : n \in \omega \rangle$. The following relationship between M and h is readily verified as the salient features are encoded into $\mathrm{css}(M_n(a))$.

$(\star)_{M,h}$: For every $n \geq 1$, $s \in Seq$, $k \in Sp(s)$, $\ell \in \omega$ such that $\mathrm{mult}(s{\restriction}_n) \geq \ell$, and $a \in M$ such that $h([a]_\infty) = s$, there are pairwise E_n-inequivalent $\{d_i : i < \ell\} \subseteq M$ such that

$$\bigwedge_{i<\ell} \Big[E_{n-1}(d_i, a) \wedge h([d_i]_\infty) = s \wedge c(d_i) = k \Big]$$

We also identify two species of elements of Seq. Call $s \in Seq$ of *isolated type* if there is $n \in \omega$ such that $\mathrm{mult}(s{\restriction}_m) = 1$ for every $m \geq n$ and of *perfect type* otherwise. The latter name is apt, as $h^{-1}(s)$ is perfect (has no isolated points) whenever s is not of isolated type. We argue that if $s \in Seq$ is of isolated type, then $Sp(s)$ is a singleton. Indeed, choose n such that $\mathrm{mult}(s{\restriction}_m) = 1$ for every $m \geq n$ and choose $a, b \in M$ such that $h([a]_\infty) = h([b]_\infty) = s$. We will show that $c(a) = c(b)$. To see this, by applying $(\star)_{M,h}$ at level $n+1$ with $k = c(b)$, get $d \in M$ such that $E_n(a, d)$, $h([d]_\infty) = s$, and $c(d) = c(b)$. But now, as $h([a]_\infty) = h([d]_\infty) = s$, the fact that s is of isolated type implies that $E_\infty(a, d)$. Thus, $c(a) = c(d) = c(b)$ as required.

We begin 'unpacking' D by inductively constructing an ω-tree $(J, \leq)$ and a surjective tree homomorphism $h : (J, \leq) \to (I, \leq)$. Begin the construction of $J = \bigcup_{n\in\omega} J_n$ by taking $J_0 = \{\rho_0\}$ to be a singleton and defining $h(\rho_0) = \sigma$. Suppose nth level J_n has been defined, together with $h : \bigcup_{j\leq n} J_j \to \bigcup_{j\leq n} I_j$. For each $\rho_n \in J_n$, we define its immediate successors $Succ_J(\rho_n)$ as follows. Look at $Succ_I(h(\rho_n)) \subseteq I_{n+1}$. For each $\sigma_{n+1} \in Succ_I(h(\rho_n))$, choose a set $A_{n+1}(\sigma_{n+1})$ of cardinality $\mathrm{mult}(\sigma_{n+1}) \in \omega + 1) \setminus \{0\}$ such that the sets $A_{n+1}(\sigma_{n+1})$ are pairwise disjoint. Let

$$Succ_J(\rho_n) := \bigcup \{ A_{n+1}(\sigma_{n+1}) : \sigma_{n+1} \in Succ_I(h(\rho_n)) \}$$

and put $J_{n+1} := \bigcup \{ Succ_J(\rho_n) : \rho_n \in J_n \}$. We extend h by $h(\rho) = \sigma_{n+1}$ for every $\rho \in A_{n+1}(\sigma_{n+1})$.

Now, having completed the construction of $(J, \leq)$ and the tree homomorphism $h : (J, \leq) \to (I, \leq)$, there is a unique extension (which we also call h) $h : [J] \to [I]$ from the branches of J to the branches of I such that $h(\eta) = s$ if and only if $h(\eta{\restriction}_n) = s{\restriction}_n$ for every $n \in \omega$.

The universe of the L-structure N we are building will be a subset of $h^{-1}(Seq) \times (\omega + 1)$ and for $(\eta, i), (\nu, j) \in N$, we will interpret E_n by

$$E_n((\eta, i), (\nu, j)) \quad \text{if and only if} \quad \eta{\restriction}_n = \nu{\restriction}_n$$

In particular, we will have $[(\eta, i)]_\infty = \{(\eta, j) : (\eta, j) \in N\}$. To finish our description of N, we must assign a 'color' to each element of $h^{-1}(Seq)$. We individually describe colors to each branch of $h^{-1}(s)$ for every $s \in Seq$. First, if s is of isolated type, then from above, we know that $Sp(s) = \{k\}$ for a single 'color' $k \leq \omega$. Accordingly, put elements $\{(\eta, i) : i < k\}$ into the universe of N for every η satisfying $h(\eta) = s$. For each $s \in Seq$ that is not of isolated type, note that $h^{-1}(s)$ has no isolated points. Thus, we can choose a partition $h^{-1}(s) = \bigcup D_k(s)$ into disjoint dense subsets indexed by colors $k \in Sp(s)$. Then, for each $\eta \in D_k(s)$ put elements $\{(\eta, i) : i < k\}$ into the universe of N. This completes our construction of the L-structure $N \in \underline{\mathrm{V}}$, and it is easily verified that this construction entails $(\star)_{N,h}$.

We now work in $\underline{\mathrm{V}}[G]$ and demonstrate that $M \equiv_{\infty,\omega} N$. Indeed, all that we need for this is that in $\underline{\mathrm{V}}[G]$, both $(\star)_{M,h}$ and $(\star)_{N,h}$ hold. Let $\mathcal{F}$ consist of all $(\overline{a}, \overline{b})$ such that $\lg(\overline{a}) = \lg(\overline{b})$, $\overline{a}$ from M, and $\overline{b}$ from N that satisfy for each $i < \lg(\overline{a})$, $c(a_i) = c(b_i)$ and $h([a_i]_\infty) = h([b_i]_\infty)$; and for each $n \in \omega$, $i < j < \lg(\overline{a})$, $M \models E_n(a_i, a_j)$ if and only if $N \models E_n(b_i, b_j)$.

To see that $\mathcal{F}$ is a back-and-forth system, choose any $(\overline{a}, \overline{b}) \in \mathcal{F}$ and choose any $a^* \in M$. We will find $b^* \in N$ such that $(\overline{a}a^*, \overline{b}b^*) \in \mathcal{F}$, and the argument in the other direction is symmetric. If $a^* \in \overline{a}$ it is obvious what to do, so assume it is not. If $E_\infty(a^*, a_i)$ for some i, then as $c(a_i) = c(b_i)$, we can find $b^* \notin \overline{b}$ such that $E_\infty(b^*, b_i)$ which suffices. Now assume that $\neg E_\infty(a^*, a_i)$ holds for each i. Let $k = c(a^*)$ and $s = h([a^*]_\infty)$.

Let $n > 0$ be least such that $\neg E_n(a^*, a_i)$ for all i. Let $A_1 = \{a_i : E_{n-1}(a^*, a_i)\}$ and let B_1 be the associated subset of $\overline{b}$. By the axioms of REF it suffices to find $b^* \in N$ such that $c(b^*) = k$, $h([b^*]_\infty) = s$, $E_{n-1}(b^*, b)$ for some/every $b \in B_1$, but $\neg E_n(b^*, b)$ for every $b \in B_1$. To find such an element, let

$$A_2 = \{a \in A_1 : \text{ there is some } a' \in [a]_n \text{ such that } c(a') = k \text{ and } h([a']_\infty) = s\}$$

Let $A_3 \subseteq A_2$ be any maximal, pairwise E_n-inequivalent subset of A_2 and let $\ell = |A_3| + 1$. The set $\{a^*\} \cup A_3$ witnesses that $\mathrm{mult}(s\restriction_n) \geq \ell$. [More precisely, for each $a \in A_3$, choose $a' \in [a]_n$ with $c(a') = k$ and $h([a']_\infty) = s$. Then $\{a^*\} \cup \{a' : a \in A_3\}$ witnesses $\mathrm{mult}(s\restriction_n) \geq \ell$.] Let B_3 be the associated subset of $\overline{b}$.

Thus, by $(\star)_{N,h}$, choose a family $\{d_i : i < \ell\}$ as there. By pigeon-hole choose an $i^* < \ell$ such that $\neg E_n(d_{i^*}, b)$ holds for all $b \in B_3$. It is easily checked that d_{i^*} is a possible choice for b^*. As noted above, this completes the proof of the Claim. □

In particular, $N \models \sigma$, establishing groundedness. □

We now turn our attention to two classical complete theories extending REF. These are often given as first examples in stability theory. We denote them by REF(inf) and REF(bin), respectively. REF(bin) is the extension of REF asserting that for every n, E_{n+1} partitions each E_n-class into two E_{n+1}-classes, while REF(inf) asserts that for all n, E_{n+1} partitions each E_n-class into infinitely many E_{n+1}-classes.

The following facts are well known.

Fact 5.2. *Both REF(bin) and REF(inf) are complete theories that admit quantifier elimination.*

- *REF(bin) is superstable but not ω-stable; and*
- *REF(inf) is stable, but not superstable.*

We will see below that these theories have extremely different countable model theory. Both are similar in that the isomorphism relation $\cong$ is not Borel. However, it turns out that REF(inf) is Borel complete, and indeed, is λ-Borel complete for every λ. That is, REF(inf) the class of models is maximally complicated, both at the countable level as well as at every uncountable level. On the other hand, REF is far from being Borel complete. In fact, its class of countable models embeds T_2 but not T_3.

5.1 Finite Branching

The subject of this subsection is REF(bin), but many of the results here easily extend to the $L_{\omega_1,\omega}$-definable class REF(fin) of 'finitely splitting, refining equivalence relations'. Recall that T_2 is a theory represented by countable sets of reals, while T_3 is a theory represented by countable sets of countable sets of reals. Our aim is to show that REF(bin) codes T_2 but not T_3, and that its isomorphism relation is not Borel. This goes as follows:

Theorem 5.3. $T_2 \leq_B$ *REF(bin).*

Proof. Begin by building a very special countable model M of REF(bin). Let S be the set of sequences from 2^ω which are eventually zero, and fix a bijection $c : S \to \underline{\mathrm{N}}$. Let M be the set of all (η, n) where $\eta \in S$ and $n < c(\eta)$. As usual, say $(\eta_1, n_1)E_m(\eta_2, n_2)$ holds if and only if η_1 and η_2 agree on the first m places. Clearly M is a model of REF(bin), and the color of (η, n) is exactly $c(\eta)$; observe that no element has color

$\aleph_0$. We will construct our models as superstructures of M, whose new elements all have color $\aleph_0$ and are not E_∞-equivalent to any element of M.

Let I be a "countable set of reals". As $2^\omega \setminus S$ is Borel isomorphic to 2^ω, we assume $I \subseteq 2^\omega \setminus S$. Let M_I be the L-structure extending M with universe $M \cup (I \times \omega)$, where again, $(\eta_1, n_1)E_m(\eta_2, n_2)$ holds if and only if $\eta_1 \restriction_m = \eta_2 \restriction_m$. To see that $M_I \ncong M_J$ when $I \neq J$, the point is that the E_∞-classes from M can be picked out from the isomorphism type of M_I: $(\eta, n)/E_\infty$ is the set of elements whose colors are precisely $c(\eta)$. So, for any $x \in M_I$ with color $\aleph_0$, define a real $r(x) \in 2^\omega$ by $r(x)(m) = 1$ if and only if $xE_m(c^{-1}(m), 0)$. That is, elements of $M_I \setminus M$ code reals in a way which is recoverable from the isomorphism type. Thus, for $I, J \subseteq (2^\omega \setminus S)$,

$$M_I \cong M_J \quad \text{if and only if} \quad \{r(\eta, 0) : \eta \in I\} = \{r(\eta, 0) : \eta \in J\} \quad \text{if and only if} \quad I = J$$

By a suitable coding, the map $I \mapsto M_I$ can be transformed into a Borel reduction $T_2 \leq_B \mathrm{REF(bin)}$. □

To show that T_3 does not embed into REF(bin), we clarify $\equiv_{\infty\omega}$-equivalence on a slightly wider class of L-structures. Let REF(fin) denote a sentence of $L_{\omega_1,\omega}$ extending REF asserting that every E_{n+1}-class partitions every E_n-class into finitely many E_{n+1}-classes.

Lemma 5.4. *Every model M of REF(fin) has an $\equiv_{\infty\omega}$-equivalent submodel $N \subseteq M$ of size at most $\beth_1$.*

Proof. For each E_∞-class $[a]_\infty \subseteq M$, let

$$B([a]_\infty) = \begin{cases} [a]_\infty & \text{if } [a]_\infty \text{ is countable} \\ \text{any countable subset of } [a]_\infty & \text{if } [a]_\infty \text{ is uncountable} \end{cases}$$

and let N be the substructure of M with universe $\bigcup\{B([a]_\infty) : a \in M\}$. It is easily seen that $N \equiv_{\infty,\omega} M$. That N has size at most continuum follows from the finite splitting at each level. □

Combined with groundedness, this gives us the nonembedding result we wanted:

Theorem 5.5. *$||\mathrm{REF(bin)}|| = ||\mathrm{REF(fin)}|| = I_{\infty,\omega}(\mathrm{REF(bin)}) = I_{\infty,\omega}(\mathrm{REF(fin)}) = \beth_2$. In particular, both REF(bin) and REF(fin) are short and $T_3 \not\leq_B$ REF(bin), REF(fin).*

Proof. Since $T_2 \leq_B \mathrm{REF(bin)}$, $\beth_2 = ||T_2|| \leq ||\mathrm{REF(bin)}||$. On the other hand, since REF is grounded, $||\mathrm{REF(fin)}|| = I_{\infty,\omega}(\mathrm{REF(fin)})$ but the latter cardinal is bounded above by $\beth_2$ by Lemma 5.4. Thus, all four cardinals are equal to $\beth_2$. So, by definition, both REF(bin) and REF(fin) are short. The nonembeddability of T_3 into either class follows from Theorem 3.10(2). □

Finally, we show that isomorphism for REF(bin) is not Borel.

Theorem 5.6. *Isomorphism on REF(bin) is not Borel.*

Proof. It is commonly known – see for example [2] – that isomorphism is Borel if and only if, for some α, $\equiv_\alpha$ is sufficient to decide isomorphism. Since $\equiv_0$ is implied by $\equiv$ and REF(bin) is a complete theory with more than one model, $\equiv_0$ does not decide isomorphism. We proceed by induction with a combined step and limit induction step. So suppose $\alpha_0 \leq \alpha_1 \leq \cdots$ are such that each $\equiv_{\alpha_n}$ does not decide isomorphism. That is, for each n, there is a pair A_n, B_n of models of REF(bin) which are nonisomorphic but where $A_n \equiv_{\alpha_n} B_n$. Let $\alpha = \sup\{\alpha_n + 1 : n \in \omega\}$. We will construct a pair A, B of models of REF(bin) where $A \equiv_\alpha B$ but $A \ncong B$. This is sufficient.

In fact we produce, for any countable $X \subset 2^\omega$, a model M_X constructed from X and the sequences $\mathcal{A} = \{A_n : n \in \omega\}$ and $\mathcal{B} = \{B_n : n \in \omega\}$. By adding a new element to each E_∞-class of each model, we may ensure the color 1 does not appear in any A_n or B_n; this will not add any dissimilarity. There are two essential claims. First: for any X and Y which are dense in 2^ω, $M_X \cong M_Y$ if and only if $X = Y$. Second: for any X and Y which are dense in 2^ω, $M_X \equiv_\alpha M_Y$. Since there are many distinct countable dense subsets of 2^ω, this proves the theorem. We now describe the construction.

For any model $N \models \text{REF(bin)}$, we can naturally "shift" the model up k levels; that is, if $a, b \in N$, then $aE_{n+k}b$ holds in the shift N^{+k} if and only if aE_nb holds in N. This is not a model of REF(bin) if $k > 0$, but is naturally an E_k class in some larger set. As before, we begin by constructing a small model M which will be contained in every M_X. We label the E_k-classes of M by elements of 2^k so that if a class refines another, the label of the first extends the label of the second. The E_2 class labeled $(0,0)$ will be exactly A_0^{+2}; the E_2-class labeled $(1,0)$ will be B_0^{+2}. Observe that this "settles" all elements whose classes' labels begin with either of these sequences. The E_2-classes $(0,1)$ and $(1,1)$ "branch," acting like the E_0-class $()$. That is, $(0,1,0,0)$ and $(1,1,0,0)$ are both A_1^{+4}, while $(0,1,1,0)$ and $(1,1,1,0)$ are both B_2^{+4}. The remaining classes continue to branch.

The resulting structure is a model of REF(bin), all of whose elements are inside some "bubble;" that is, a shifted copy of some A_n or B_n, and do not have color 1. We expand it to M_X, given a countable $X \subset 2^\omega$, by adding new elements of color 1 which are not in any bubble. Specifically, for each $\eta \in X$, we add a new x_η to M whose E_{2n}-classes' labels all end in 1 (that is, x_η does not fall into a bubble) and whose E_{2n+1}-classes' label always end in $\eta(n)$. Refer to these new elements as "branches," so that M_X consists of branches and (elements of) bubbles. We can now prove our first claim:

Claim 1: Let $X, Y \subset 2^\omega$ be countable and dense. Then $M_X \cong M_Y$ if and only if $X = Y$.
Proof: If $X = Y$ then $M_X = M_Y$ setwise, so they are isomorphic. On the other hand, observe that since X is dense, an E_{2n}-class in M_X is a bubble if and only if it does not contain an element of color 1. Further, the labeling of the E_1-classes can be recovered – the class labeled (0) is the one containing a bubble at level 2 which is isomorphic to A_1^{+2}. Consequently we can fully label the E_2-classes as well, and by repeating this idea, recover all the labeling everywhere. The same logic holds for M_Y. All this is preserved under isomorphism, so if $M_X \cong M_Y$, the isomorphism takes bubbles to bubbles and thus preserves all labeling. It also takes color 1 elements to color 1 elements, so takes $\eta \in X$ to $\eta \in Y$. Thus, $X = Y$. □

The second claim is somewhat more involved. We actually expand the language somewhat, noting that if two structures are $\equiv_\alpha$-equivalent in a larger language, they're $\equiv_\alpha$-equivalent in the smaller language as well. So to our structures we add unary predicates for "being in a bubble which begins at level $2n$." We do not add predicates for the labeling of the structure, as the high level of similarity comes from the difficulty in determining what that labeling is.

Claim 2: Let $X, Y \subset 2^\omega$ be countable and dense. Then $M_X \equiv_\alpha M_Y$.
Proof: It is sufficient to show that for all $\beta < \alpha$, $M_X \equiv_{\beta+1} M_Y$. Once a particular β has been chosen, there is a (least) N where $\alpha_n \geq \beta$ for all $n \geq N$. For each $m \leq N$, pick an element $a_m \in A_m$ and $b_m \in B_m$, and add constants to M_X and M_Y for each (shifted copy of) a_m and b_m in each model; these effectively label the bubbles of height at most $2N$. We will show that $M_X \equiv_{\beta+1} M_Y$ in the expanded language. Our goal is to define a strategy for player II that is well-defined for all possible plays (where II follows her strategy) and which preserves the atoms in the expanded language. In fact it will also satisfy that branches go to branches, that bubbles go to bubbles at the same level, that between (shifts of) A_n and A_n or B_n and B_n, the identity strategy is followed, and between (shifts of) A_n and B_n, a fixed strategy for the $\equiv_\beta$-game is followed.

For the first play of the game, if I plays an element of a bubble in M_X or M_Y, II plays the same element of the same bubble in the other model. If I plays a branch, let II play a branch in the other model which agrees with the labeling up to level $2N$. Both plays are possible since X and Y are dense, and the play preserves every atom in the new language as well as our inductive hypothesis above.

Now suppose our play of length $n+1$ consists of $\emptyset a$ versus $\emptyset b$, along with the chosen ordinals $\beta = \beta_0 > \cdots > \beta_{n-1}$. Notationally, we consider the constants for elements of M_X to be negative-indexed elements of $\emptyset a$, and similarly with M_Y and $\emptyset b$. Then I plays $\beta_n < \beta_{n-1}$ and a_n. We have four cases. First and most trivially, if $a_n = a_i$ for some $i < n$, let $b_n = b_i$. Clearly this is legal and preserves the inductive hypothesis.

Second, suppose a_n is from a bubble A from M_X which contains some a_i with $i < n$. Then there is a corresponding bubble B in M_Y containing b_i. Let $\emptyset c \subset \emptyset a$ be the portion of the play which is in A, and $\emptyset d \subset \emptyset b$ which is in B. Also let $\overline{\gamma}$ be the ordinals played by I alongside $\emptyset c$. By induction, there is a strategy between A and B, which the play $(\emptyset c, \emptyset\gamma)$ and $\emptyset d$ follows. Consequently, there is a response b_n where $(\emptyset c a_n, \emptyset\gamma\beta_n)$ and $\emptyset d b_n$ follow the strategy. This satisfies the inductive hypothesis, and by transitivity

of the E_i, play satisfies all the atoms of the expanded language. The cases where $a_n \in M_Y$, or where A contains a b_i instead of an a_i, are identical up to notation.

Third, if a_n is a branch, then it must be from a new branch, since all branches have color 1. There is a maximal m where (for some $i < n$) either $a_n E_m a_i$ or $a_n E_m b_i$. We assume the former; the latter is identical up to notation. Since the branches are dense (outside the bubbles), there is a branch b_n in the other model where $b_n E_m b_i$ but not $b_n E_{m+1} b_i$. Choose any one of them as our play; there are no definability concerns, so use the axiom of choice. This is well-defined by density of X and Y, preserves all the atoms of the expanded language, and satisfies the inductive hypothesis.

Fourth and finally, say a_n is from a bubble A from M_X (or M_Y) where for all $i < n$, $a_i, b_i \notin A$. Let n be the height of the beginning of the bubble A, observing that $n > 2N$ by virtue of our original expansion of the language. Let m be maximal where some element of A E_m-agrees with some a_i or b_i in M_X, noting that $m \leq n$. Let B be a bubble of height n from M_Y (respectively, M_X) where B E_m-agrees with the corresponding b_i or a_i in M_Y but does not E_{m+1}-agree with it. Then B may or may not be a bubble of the same type as A, but there is a strategy between A and B which works up to $\equiv_\beta$ since $n > 2N$. Let b_n be the result of following this strategy with first play (a_n, β_n), noting that $\beta_n < \beta$, so this is legal. This satisfies all atoms and fulfills the inductive hypothesis. These cases are exhaustive, so the strategy is demonstrated and the proof is complete. □

With both claims finished, let $X \subset 2^\omega$ be the set of sequences which are eventually zero, and $Y \subset 2^\omega$ be the set of sequences which are eventually one. Then both are countable, so M_X and M_Y are well-defined. Both are dense, so $M_X \equiv_\alpha M_Y$. And $X \neq Y$, so $M_X \ncong M_Y$. This completes the induction and the proof. □

This gives the first known example of the following behavior:

Corollary 5.7. *There is a complete first-order theory for whom isomorphism is neither Borel nor Borel complete.*

Here, the example is REF(bin), the paradigmatic example of a superstable, non-ω-stable theory. Thus we might informally expect this behavior to be extremely common for such theories. Since isomorphism is not Borel, we cannot truly consider REF(bin) to be especially tame. However, the theory is relatively simple in the sense that it cannot code much infinitary behavior. We end with the following class of examples which follow naturally from this one:

Corollary 5.8. *For any α with $2 \leq \alpha < \omega_1$, there is a complete first-order theory S_α whose isomorphism relation is not Borel, and where $T_\beta \leq_B S_\alpha$ if and only if $\beta \leq \alpha$.*

Each of these theories is grounded, superstable, but not ω-stable.

Proof. Take $S_2 :=$ REF(bin). We construct $S_{\alpha+1}$ as $J(S_\alpha)$, and for limit α, construct S_α as $\prod_\beta S_\beta$ as β varies below α. That $T_\alpha \leq_B S_\alpha$ follows from induction and Propositions 3.17 and 3.19, part (4). That $T_{\alpha+1} \not\leq_B S_\alpha$ follows from the fact that $\|S_\alpha\| = \|T_\alpha\| < \|T_{\alpha+1}\|$, which follows from part (3) of those Propositions.

Groundedness follows from part (2), and the place in the stability spectrum is a standard type-counting argument, beginning with the fact that REF(bin) has the desired properties. □

We end this subsection with an open question:

Question 5.9. Let α be 0 or 1. Is there a first-order theory S_α whose isomorphism relation is not Borel, and where $T_\beta \leq_B S_\alpha$ if and only if $\beta \leq \alpha$?

Note that the instance of the above question for $\alpha = 0$ is precisely Vaught's conjecture for first-order theories. The instance for $\alpha = 1$ is related to so-called Glimm-Effros dichotomies, which have been shown in [5] to hold for a number of nice classes of models.

5.2 Infinite Branching

We now turn our attention to REF(inf) specifically. We have really only one major theorem:

Theorem 5.10. *REF(inf) is Borel complete. Indeed, for each infinite cardinal λ, REF(inf) is λ-Borel complete.*

Proof. By Theorem 3.11 of [12], it is enough to produce a λ-Borel reduction from subtrees of $\lambda^{<\omega}$, up to $\equiv_{\infty\omega}$, to models of REF(inf), also up to $\equiv_{\infty\omega}$. The models produced will have size λ. Let $I \subset \lambda^\omega$ be the set of all sequences which are eventually zero.

Now, let $S \subset \lambda^{<\omega}$ be a subtree, that is, a downward closed subset. The E_∞-classes of M_S will be $(e_\eta : \eta \in I)$, so that $e_\eta E_n e_\nu$ if and only if $\eta{\restriction_n} = \nu{\restriction_n}$. Let e_η have color two if and only if, for some $t \in \lambda^{<\omega} \setminus S$, $t \frown (0) \subset \eta$. Otherwise, e_η will have color one. Consequently, we can easily recover the *complement* of S, up to tree isomorphism, by saying $t \notin S$ if and only if every class above $t \frown (0)$ has color two. Note the reason for the (0) is so that if $t \in S$, there is at least one branch – t followed by infinite zeros – which has color one. If t is a terminal node in S, this will be the only such branch.

It is then an easy exercise to show that $S \equiv_{\infty\omega} S'$ if and only if $M_S \equiv_{\infty\omega} M_{S'}$, which is as desired. Note that in the case $\lambda = \aleph_0$, this is a direct Borel reduction from subtrees of $\omega^{<\omega}$ to $(\mathrm{Mod}(\mathrm{REF(inf)}), \cong)$, giving Borel completeness. □

Corollary 5.11. *REF(inf) is not short. Indeed, REF(inf) has class-many $\equiv_{\infty\omega}$-inequivalent models in $\underline{V}$.*

6 ω-Stable Examples

Here we discuss two more first-order theories whose isomorphism relations are not Borel, but where one is Borel complete, and the other does not embed T_3. Interestingly, both are extremely similar model-theoretically. Both are ω-stable with quantifier elimination, and have ENI-NDOP and eni-depth 2, which together give a strong structure theorem in terms of stability theory. Indeed, as we will show, both have exactly $\beth_2$ models up to $\equiv_{\infty\omega}$, meaning that at a "macro" level, they are extremely similar. Yet at a "micro" level, they are quite different. Because of the similarity of the examples, we will be able to highlight exactly why one is relatively simple, and the other is not.

So let us define the theories. The first, K, is due to Koerwien and constructed in [11]. The language has unary sorts U, V_i, and C_i, as well as unary functions S_i and π_i^j for $i \in \omega$ and $j \leq i+1$. The axioms are as follows:

- The sorts U, V_i, and C_i are all disjoint. U and each of the V_i are infinite, but each C_i has size 2.
- π_i^{i+1} is a function from V_i to U; π_i^j is a function from V_i to C_j when $j \leq i$.
- For each tuple $\o c = (c_0, \ldots, c_i)$ and each $u \in U$, $\pi_i^{-1}(\o c, u)$ is nonempty. Here π_i refers to the product map $\pi_i^0 \times \cdots \times \pi_i^{i+1} : V_i \to C_0 \times \cdots C_i \times U$.
- S_i is a unary successor function from V_i to itself, and $\pi_i \circ S_i = \pi_i$.

We have a few remarks. Typically we will drop the subscript on π_i if it is clear from context. There is a slight ambiguity about the sorts, whether one works in traditional first-order logic (and thus there may be "unsorted" elements) or in multisorted logic (where there will not be). Our position is that the unsorted elements exist, but are unimportant for what we're doing here. Therefore we will often ignore them.

The properties of K have been well studied by Koerwien in [11]; we summarize his findings here:

Theorem 6.1. *K is complete with quantifier elimination. It is ω-stable, so has constructible models over arbitrary sets, has ENI-NDOP, and is eni-shallow of eni-depth 2. Furthermore, the isomorphism relation for K is not Borel.*

Our other theory is a tweak of K, so we call it TK. The language is slightly different; we have unary sorts U, V_i, and C_i as before, but have unary functions S_i, π_i^0, π_i^1, and τ_{i+1} for $i \in \omega$. The axioms are as follows:

- The sorts U, V_i, and C_i are all disjoint. U and each of the V_i are infinite, but each C_i has size 2^i.
- τ_{i+1} is a surjection from C_{i+1} to C_i where, for all $c \in C_i$, $|\tau_{i+1}^{-1}(c)| = 2$.
- π_i^1 is a function from V_i to U; π_i^0 is a function from V_i to C_i.
- For each tuple $c \in C_i$ and each $u \in U$, $\pi_i^{-1}(c, u)$ is nonempty. Here π_i refers to the product map $\pi_i^0 \times \pi_i^1 : V_i \to C_i \times U$.
- S_i is a unary successor function from V_i to itself, and $\pi_i \circ S_i = \pi_i$.

The preceding notes also apply to K. The behavior is also extremely similar, and essentially the same proofs of basic properties of K apply to TK. We summarize this now:

Theorem 6.2. *TK is complete with quantifier elimination. It is ω-stable, so has constructible models over arbitrary sets, has ENI-NDOP, and is eni-shallow of eni-depth 2.*

We can easily see that both K and TK have relatively few models up to back-and-forth equivalence:

Proposition 6.3. $I_{\infty,\omega}(K) = I_{\infty,\omega}(TK) = \beth_2$.
Indeed, every model M of either theory has a submodel N where $M \equiv_{\infty\omega} N$ and $|N| \leq \beth_1$.

Proof. Let T be either K or TK. For the proof of the proposition we can restrict attention to models of T with a fixed algebraic closure of the emptyset $\bigcup_i C_i$. Let C be either all finite sequences $(a_j : j < i)$, where each $a_j \in C_j$ (if $T = \text{K}$) or else $\bigcup_i C_i$ (if $T = \text{TK}$).

We first show $I_{\infty,\omega}(T) \geq \beth_2$. For each $\eta \in 2^\omega$, let u_η be some element which will eventually be part of U in some model of T. For any $n \in \omega$ and any $c \in C$ where $\pi_n^{-1}(u, c)$ is defined, we insist the S-dimension of $\pi_n^{-1}(u, c)$ be infinite if $\eta(n) = 1$, or equal to one otherwise. For any infinite $X \subseteq 2^\omega$, define M_X to have $U^{M_X} = \{u_\eta : \eta \in X\}$ with the described behavior of the V_i and S_i. Evidently if $Y \subseteq 2^\omega$ is infinite and $X \neq Y$, then for any $\eta \in X \setminus Y$, there is no $\nu \in Y$ where $(M_X, u_\eta) \equiv_{\infty\omega} (M_Y, u_\nu)$, and symmetrically. Thus, $M_X \not\equiv_{\infty\omega} M_Y$. Since there are $\beth_2$ infinite subsets of 2^ω, $I_{\infty\omega}(T) \geq \beth_2$.

That $I_{\infty\omega}(T) \leq \beth_2$ follows immediately from the second claim. So let M be some model of T, of any particular cardinality. We begin by stripping down the V_i. For every $u \in U$ and $c \in C$, if $\pi^{-1}(c, u)$ is uncountable, drop all but a countable S_i-closed subset of infinite S_i-dimension. Do this for all pairs (c, u). The result is $\equiv_{\infty\omega}$-equivalent to the original by an easy argument, and $\pi^{-1}(c, u)$ is now always countable.

Next we need only ensure that U has size at most continuum. So put an equivalence relation E on U, where we say uEu' holds if and only if, for all $c \in C$, the dimensions of $\pi^{-1}(c, u)$ and $\pi^{-1}(c, u')$ are equal. If any E-class is uncountable, drop all but a countable subset; the resulting structure is $\equiv_{\infty\omega}$-equivalent to the original again. Further, each E-class is now countable, and there are only $|C^\omega| = \beth_1$ possible E-classes, so the structure now has size at most $\beth_1$. This completes the proof. □

Any additional complexity of either theory comes from elementary permutations of the algebraic closure of the empty set. In any model M of either K or TK, $\text{acl}_M(\emptyset) = \bigcup_{i \in \omega} C_i(M)$. In models M of TK, the projection functions $\{\tau_i\}$ naturally induce a tree structure, so we think of $\text{acl}_M(\emptyset)$ as being a copy of $(2^{<\omega}, \leq)$. In models M of K, as each $C_i(M)$ has exactly two elements, one can naïvely think of $\text{acl}_M(\emptyset)$ as being indexed by $2 \times \omega$. However, in order to better understand the action of functions π_j^i, it is better to think of $\text{acl}_M(\emptyset)$ as consisting of finite sequences $\langle a_j : j < i \rangle$, where each $a_j \in C_j(M)$. These finite sequences, when ordered by initial segment, also give a natural correspondence of $\text{acl}_M(\emptyset)$ with the tree $(2^{<\omega}, \leq)$. Henceforth, when discussing models M of either K or TK, we will view $\text{acl}_M(\emptyset)$ as being indexed by the tree $(2^{<\omega}, \leq)$.

Next, we discuss the group G of elementary permutations of $\text{acl}_M(\emptyset)$ in a model of one of the two theories. For K, the relevant group is $(2^\omega, \oplus)$, the direct product of ω copies of the two-element group. Indeed, in

any model of K, any sequence of permutations of elements of the sets $C_i(M)$ is an elementary permutation of $\mathrm{acl}_M(\emptyset)$. Moreover, any $g \in G$ naturally induces an automorphism of the induced tree $(2^{<\omega}, \leq)$. In TK, as elementary permutations have to respect the τ_i structure, the relevant group of elementary permutations is $\mathrm{Aut}(2^{<\omega}, \leq)$. For both theories, each $g \in G$ induces a continuous action on $\mathrm{Mod}(T)$. We see in the proposition below that for both theories, $(\mathrm{Mod}(T), \cong)$ is Borel equivalent to $([\mathcal{N}]^{\aleph_0}, E_G)$, where E_G is the orbit equivalence relation of a natural, continuous action of G on $\mathcal{N}$. Both groups of elementary permutations of $\mathrm{acl}(\emptyset)$ are compact, but only the first is abelian. It turns out that being abelian is enough to produce relative simplicity, while being nonabelian leaves enough room to allow TK to be Borel complete.

Proposition 6.4. *Let T be either K or TK, and let G be either $(2^\omega, \oplus)$ or $Aut(2^{<\omega}, \leq)$, respectively. Let $\mathcal{N}$ be Baire space ω^ω. Then is a strongly definable, persistent action of G on X such that:*

- $(Mod(T), \cong) \sim_B (\mathcal{N}^\omega, E_{G \times S_\infty})$.
- $(Mod(T), \cong)$ *and* $(\mathcal{P}_{\aleph_1}(\mathcal{N}), E_G)$ *are $\leq_{HC}$-biembeddable.*

Here E_G is the orbit equivalence relation of G on $\mathcal{N}$, as well as of the diagonal action of G on $\mathcal{P}_{\aleph_1}(\mathcal{N})$. G acts diagonally on sequences from X^ω and this commutes with the permutation action of S_∞ on $\mathcal{N}^\omega$. So $G \times S_\infty$ acts naturally on $\mathcal{N}^\omega$; $E_{G \times S_\infty}$ is the orbit equivalence relation of this action.

Proof. Fix a bijection $f : (2^{<\omega} \setminus \{\emptyset\}) \to \omega$. Using f we can conflate the strongly definable Polish spaces $(\omega + 1 \setminus \{\emptyset\})^{2^{<\omega} \setminus \{\emptyset\}}$ and $\mathcal{N} = \omega^\omega$. Also, G acts on $2^{<\omega} \setminus \{\emptyset\}$ naturally; if $G = (2^\omega, \oplus)$, then $g \cdot \sigma = g\restriction_{|\sigma|} \oplus \sigma$. If $G = \mathrm{Aut}(2^{<\omega}, \leq)$, then $g \cdot \sigma = g(\sigma)$. Certainly both are strongly definable, persistent actions.

To show $(\mathrm{Mod}(T), \cong) \leq_B (\mathcal{N}^\omega, E_{G \times S_\infty})$, first let $M \in \mathrm{Mod}(T)$ be arbitrary. As described in the discussion preceding this proposition, we may identify $\mathrm{acl}_M(\emptyset)$ with $2^{<\omega}$ in a Borel way, and U^M with ω, using the original indexing of the universe of M with ω. Then each element $u \in U^M$ codes a function $c_u : (2^{<\omega} \setminus \{\emptyset\}) \to (\omega + 1 \setminus \{\emptyset\})$, where $c_u(\sigma)$ is the $S_{|\sigma|}$-dimension of $\pi^{-1}(u, \sigma)$. Under the bijections described above, each c_u can be considered as an element x_u of X. So take M to the sequence $S_M = (x_n^M : n \in \omega)$, where $x_n = x_u$ for the unique $u \in U^M$ corresponding to the the number n.

Suppose $N \in \mathrm{Mod}(T)$ as well; again, we may assume $U^N = \omega$ and $\mathrm{acl}_N(\emptyset) = 2^{<\omega} \setminus \emptyset$. If $\sigma : M \cong N$, then σ induces a $g \in G$ which takes $\mathrm{acl}_M(\emptyset)$ to $\mathrm{acl}_N(\emptyset)$ and a permutation $\tau \in S_\infty$ taking U^M to U^N where for all $u \in U^M$, $c_{\sigma u} = g(c_u)$. Consequently $(g, \sigma)(x_n^M) = x_n^N$ for all $n \in \omega$, so $S_M E_{G \times S_\infty} S_N$. The converse is similar, so $M \mapsto S_M$ is a Borel reduction.

To show $(\mathcal{N}^\omega, E_{G \times S_\infty}) \leq_B (\mathrm{Mod}(T), \cong)$, fix a sequence $S = (x_n : n \in \omega)$. We describe the case for the theory K, as the presence of the functions $\{\tau_i\}$ make the case of TK easier. By the identifications from the beginning, we may assume each x_n is a function from $(2^{<\omega} \setminus \{\emptyset\})$ to $(\omega + 1 \setminus \{\emptyset\})$. We define M_S to have $U^{M_S} = \omega$, and have $C_i^{M_S} = 2$. Then we put the $S_{|\sigma|}$-dimension of $\pi^{-1}_{|\sigma|}(x_n, \sigma)$ to be $x_n(\sigma)$; that is, we use x_n as a function to determine what its various S-dimensions should be. Then $S_1 E_{G \times S_\infty} S_2$ if and only if $M_{S_1} \cong M_{S_2}$, as desired; this completes the proof of the Borel portion.

Showing $(\mathcal{N}^\omega, E_{G \times S_\infty}) \leq_{\mathrm{HC}} (\mathcal{P}_{\aleph_1}(X), E_G)$ is almost trivial; we only need to eliminate multiplicity concerns. So fix $S \in \mathcal{N}^\omega$, and for each $x \in S$, define $m_S(x)$ to be $|\{y \in S : x E_G y\}| + 1$ if finite, or zero if infinite. This way $m_S(x)$ codes the multiplicity of x by a natural number. Then let $p(x) \in \mathcal{N}$ take $n \in \omega$ to the pair $\langle x(n), m(x) \rangle$, and let $p(S) = (p(x_n) : n \in \omega)$. The composition $S \mapsto p(S)$ is a Borel reduction: $S_1 E_{G \times S_\infty} S_2$ if and only if $p(S_1)\, E_{G \times S_\infty}\, p(S_2)$.

Now define the function $r : \mathcal{N}^\omega \to \mathcal{P}_{\aleph_1}(\mathcal{N})$ where if $S = (x_n : n \in \omega)$ is in $\mathcal{N}^\omega$, then $r(S) = \{x_n : n \in \omega\}$; this is strongly definable and just drops the multiplicity information. Since the multiplicity is already coded by the function p, the composition $S \mapsto r(p(S))$ shows $(\mathcal{N}^\omega, E_{G \times S_\infty}) \leq_{\mathrm{HC}} (\mathcal{P}_{\aleph_1}(\mathcal{N}), E_G)$, as desired.

The reverse embedding $(\mathcal{P}_{\aleph_1}(X), E_G) \leq_{\mathrm{HC}} (\mathrm{Mod}(T), \cong)$ is similar to the preceding. □

6.1 Koerwien's Example

For this subsection, we want to prove that K is not Borel complete, and further, to characterize exactly which T_α embed in K. Recall that T_2 is a first-order theory represented by countable sets of reals, T_3 is a first-order

theory represented by countable sets of countable sets of reals, and so on. To do this, we show directly that $T_2 \leq_B \mathrm{K}$, and then show that $\mathrm{CSS(K)_{ptl}}$ has size exactly $\beth_2$. The former is quite straightforward:

Proposition 6.5. *$T_2 \leq_B K$.*

Proof. Let $X \subset 2^\omega$ be countable; we describe a model $M_X \models \mathrm{K}$ from which X can be easily recovered. Let U be the set $A \cup X$, where A is some countable infinite set which is disjoint from X. Let each C_i be the set $\{0, 1\}$ with appropriate tagging. For each tuple $(a, \varnothing c)$ with $a \in A$, give it the color "five." For each tuple $(x, \varnothing c)$ with $x \in X$, give it the color $x(|\varnothing c|)$, where we view x as a function $\omega \to 2$ and "color zero" means infinite dimension. Clearly, if $X = Y$, then $M_X \cong M_Y$. If $M_X \cong M_Y$, and $x \in X$, then there is an element $y \in U^{M_Y}$ whose colors are all zero and one, so that $y \in Y$. Further, $x(n) = 0$ if and only if y has color zero on tuples of length n, if and only if $y(n) = 0$. So $x = y$, and so $x \in Y$. The converse holds as well, so $X = Y$ if and only if $M_X \cong M_Y$.

To make this completely precise, in fact we take X from the space of sequences $(2^\omega)^\omega$. The function $X \mapsto M_X$ is a Borel function from $(2^\omega)^\omega$ to Mod(K), and it is clear that if X and Y are sequences, then $M_X \cong M_Y$ if and only if the sets $\{x_n : n \in \omega\}$ and $\{y_n : n \in \omega\}$ are equal, so the function $X \mapsto M_X$ is a Borel reduction $T_2 \leq_B \mathrm{K}$. □

Having accomplished this, we can state everything we need about K:

Theorem 6.6. *$||K|| = \beth_2$. Therefore, K is not Borel complete; indeed, there is no Borel embedding of T_3 into Mod(K).*

Proof. That $||K|| \geq \beth_2$ follows immediately from Proposition 6.3. Since $G = (2^\omega, \oplus)$ is compact and abelian, taking X to be Baire space, $||(\mathcal{P}_{\aleph_1}(X), E_G)|| \leq \beth_2$ by Theorem 4.7. As $||\mathrm{K}|| = ||(\mathcal{P}_{\aleph_1}(X), E_G)||$ by Proposition 6.4, we conclude that $||K|| = \beth_2$. That there is no Borel embedding of T_3 into Mod(K) is immediate from Theorem 3.10(2) and Corollary 3.21. □

Once we have one such example, we can apply the usual constructions to get a large class of ω-stable examples:

Corollary 6.7. *For each ordinal α with $2 \leq \alpha < \omega_1$, there is an ω-stable theory S_α whose isomorphism relation is not Borel and where $T_\beta \leq_B S_\alpha$ if and only if $\beta \leq \alpha$.*

Proof. Let S_2 be K. Then proceed inductively as in Corollary 5.8. □

There is no such example when $\alpha = 0$, since Vaught's Conjecture holds for ω-stable theories. It is unknown if there is an example when $\alpha = 1$, although it would have to be different from the existing examples. This is because for each α, our T_α have eni-depth α. However, ω-stable theories with eni-depth 1 have smooth (and thus Borel) isomorphism relation [10].

6.2 A New ω-Stable Theory

We now consider TK, with the aim of showing it is Borel complete. Indeed we have already shown $(\mathrm{Mod(TK)}, \cong)$ is Borel equivalent to $(\mathcal{N}^\omega, E_{G \times S_\infty})$, so we need only show the latter is Borel complete. Since we want to work with the group action directly, it will be more concrete to work with the original Polish space $X = \omega^{(2^{<\omega})}$, the space of colorings of $2^{<\omega}$, rather than identifying this space with $\mathcal{N}$. Recall that $G = \mathrm{Aut}(2^{<\omega}, \leq)$ acts on X by permuting the fibers; that is, for any $c : 2^{<\omega} \to \omega$, any $g \in G$, and any $\eta \in 2^{<\omega}$, $(g \cdot c)(\eta) = c(g \cdot \eta)$. Then G acts on X^ω diagonally, while S_∞ acts on X^ω by permuting the fibers, so these actions commute with one another and induce an action of the product group $G \times S_\infty$. Proposition 6.4 states, among other things, that $(X^\omega, E_{G \times S_\infty}) \leq_B (\mathrm{Mod(TK)}, \cong)$; thus, to show TK is Borel complete, it is enough to show $(X^\omega, E_{G \times S_\infty})$ is Borel complete, which we do directly.

Theorem 6.8. *$(Graphs, \cong) \leq_B (X^\omega, E_{G \times S_\infty})$.*

Proof. To simplify notation, for the whole of this proof we write E in place of $E_{G\times S_\infty}$. We need some setup first. Observe that G naturally acts on 2^ω, the set of branches of $(2^{<\omega}, \leq)$, by $g\cdot\sigma = \bigcup_n g\cdot\sigma{\restriction_n}$; this is a well-defined sequence precisely because g is a tree automorphism. Let $\{D_i : i\in\omega\}$ be a countable set of mutually dense, disjoint, countable subsets of 2^ω, and let $D = \bigcup_i D_i$. We need one claim, where we use the relative complexity of G (it would not go through if we replaced TK with K):

Claim: For any $\sigma\in S_\infty$, there is a $g\in G$ where for all $i\in\omega$, $g\cdot D_i = D_{\sigma(i)}$ as sets.
Proof: We construct such a g by a back-and-forth argument. So let $\mathcal{F}$ be the set of finite partial functions from 2^ω to itself, satisfying all the following:

- For each $f\in\mathcal{F}$ and each $\eta\in 2^\omega$, if $f(\eta)$ is defined and $\eta\in D_i$, then $f(\eta)\in D_{\sigma(i)}$.
- For each $f\in\mathcal{F}$ and each $\eta,\nu\in 2^\omega$, if both $f(\eta)$ and $f(\nu)$ are defined, then $\lg(\eta\wedge\nu) = \lg(f(\eta)\wedge f(\nu))$, where $\eta\wedge\nu$ denotes the longest common initial segment of η and ν.
- The previous conditions, but with f^{-1} and σ^{-1} instead of f and σ.

Once we establish that $\mathcal{F}$ is a back-and-forth system, then $\mathcal{F}$ defines a $g\in G$ with the desired property. For pick a branch through $\mathcal{F}$ whose union f includes all of D in both domain and range. If $s\in 2^n$, let $g(s)$ be $f(\eta){\restriction_n}$ for any η extending s; because of the consistency properties of $\mathcal{F}$, this is well-defined. Because D is dense, g defines a group automorphism, and clearly has the desired property with respect to σ.

So we need only show that $\mathcal{F}$ is a back-and-forth system. Of course the empty function is in $\mathcal{F}$. So say $f\in\mathcal{F}$ and $\eta\in 2^\omega$; we want $f'\supset f$ in $\mathcal{F}$ with $\eta\in\mathrm{dom}(f')$. We assume $\eta\notin\mathrm{dom}(f)$ already, or else this is trivial. Let n be maximal among $\{\lg(\eta\wedge\nu) : \nu\in\mathrm{dom}(f)\}$, and let $\nu\in\mathrm{dom}(f)$ be such that $\lg(\eta\wedge\nu) = n$. We then pick an element $f(\eta)$ of 2^ω which extends $f(\nu){\restriction_n}\frown(1-f(\nu)(n))$. That is, $f(\eta)$ agrees with $f(\nu)$ before stage n, but disagrees with it at n. If $\eta\in D_i$ for some i, choose this element from $D_{\sigma(i)}$, which is possible by density. If $\eta\notin D$, choose this element not to be in D, which is possible since D is countable and thus codense. This clearly satisfies the desired properties, and the other direction is symmetric, proving the claim. □

Given $\eta,\tau\in 2^\omega$ and $n>0$, let $c^n_{\eta,\tau} : 2^{<\omega}\to\omega$ be the coloring which sends $s\in 2^{<\omega}$ to n, if $s\subset\eta$ or $s\subset\tau$, or 0 otherwise. For clarity, observe that $c^n_{\eta,\tau}$ is n on a closed subset of $2^{<\omega}$ and 0 elsewhere. Finally, fix a bijection $\rho:\omega\to D\times D$. We have now fixed enough notation and can describe our coding.

Let R be a graph on ω – that is, R is a binary relation on ω which is symmetric and irreflexive. For each $n\in\omega$, $\rho(n)$ is a pair $(\eta,\tau)\in D_i\times D_j$ for some i and j. If $i=j$, define $c_n = c^1_{\eta,\tau}$. If $i\neq j$ and $\{i,j\}\in R$, then let $c_n = c^2_{\eta,\tau}$. Otherwise let $c_n = c^3_{\eta,\tau}$. Then $f(R) = \emptyset c = (c_n : n\in\omega)$ is an element of X^ω, constructed in a Borel fashion.

Suppose $\sigma : R\to R'$ is a graph isomorphism. By the claim, there is a $g\in G$ where for all $i\in\omega$, $g\cdot D_i = D_{\sigma(i)}$. It is not difficult to see then that $f(R)Ef(R')$. Indeed, suppose $c^1_{\eta,\tau}$ is some element in the sequence $f(R)$. Then for some i, $\eta,\tau\in D_i$ and so $g(\eta),g(\tau)\in D_{\sigma(i)}$. Then $g\cdot c^1_{\eta,\tau} = c^1_{g(\eta),g(\tau)}$ is some element of the sequence $f(R')$. Similarly if $c^2_{\eta,\tau}\in f(R)$, there is some $i\neq j$ where $\eta\in D_i$, $\tau\in D_j$, and $\{i,j\}\in R$. Since $\sigma : R\to R'$ is a graph isomorphism, $\{\sigma(i),\sigma(j)\}\in R'$, so if $\eta'\in D_{\sigma(i)}$ and $\tau'\in D_{\sigma(j)}$, then $c^2_{\eta',\tau'}\in f(R')$. In particular, $g\cdot c^2_{\eta,\tau}$ is in the sequence $f(R')$. Then case $c^3_{\eta,\tau}$ is similar, as is the reverse. Thus, some permutation of $g\cdot f(R)$ is equal to $f(R')$, so $f(R)Ef(R')$.

It only remains to show that if $f(R)Ef(R')$, then $R\cong R'$. We accomplish this by recovering enough of R from $f(R)$. To suppose $\emptyset c = (c_0,c_1,\ldots)\in X^\omega$. We make three definitions:

- Let $I_{\emptyset c}$ be the set of all $\eta\in D$ where for some $n<\omega$, $c_n = c^1_{\eta,\eta}$.
- Let $E_{\emptyset c}\subset I_{\emptyset c}\times I_{\emptyset c}$ be the set of all $(\eta,\tau)\in D^2$ where for some n, $c_n = c^1_{\eta,\tau}$.
- Let $R_{\emptyset c}\subset I_{\emptyset c}\times I_{\emptyset c}$ be the set of all $(\eta,\tau)\in D^2$ where for some n, $c_n = c^2_{\eta,\tau}$.

The intention is for $I_{\emptyset c}$ to be the "elements" of the graph, before modding out by equality. $E_{\emptyset c}$ is the equivalence relation given by equality. R is the graph relation for the graph. For any $\emptyset c\in X^\omega$ and $\sigma\in S_\infty$,

it is easy to check that $I_{\varnothing c} = I_{\sigma\varnothing c}$, $E_{\varnothing c} = E_{\sigma\varnothing c}$, and $R_{\varnothing c} = R_{\sigma\varnothing c}$. Similarly, for any $g \in G$, $g \cdot I_{\varnothing c} = I_{g\cdot\varnothing c}$, $g \cdot E_{\varnothing c} = E_{g\cdot\varnothing c}$, and $g \cdot R_{\varnothing c} = R_{g\cdot\varnothing c}$.

Let $\mathcal{A} \subset X^\omega$ be the set of all $\varnothing c \in X^\omega$ where $E_{\varnothing c}$ is an equivalence relation on $I_{\varnothing c}$ with infinitely many classes, and where $R_{\varnothing c}$ induces a graph on $I_{\varnothing c}/E_{\varnothing c}$; that is, R is well-defined on E-classes and is symmetric and irreflexive. By the preceding, if $\varnothing c, \varnothing d \in \mathcal{A}$ and $\varnothing c E \varnothing d$, then $(I_{\varnothing c}/E_{\varnothing c}, R_{\varnothing c}) \cong (I_{\varnothing d}/E_{\varnothing d}, R_{\varnothing d})$, where the isomorphism is given precisely by the group element taking $\varnothing c$ to $\varnothing d$.

Suppose R is a graph on ω. Then $\varnothing c = f(R) \in \mathcal{A}$, $I_{\varnothing c} = D$, the $E_{\varnothing c}$-classes are precisely the D_i, and $D_i R_{\varnothing c} D_j$ if and only if iRj. That is, $R \cong (I_{\varnothing c}/E_{\varnothing c}, R_{\varnothing c})$. So clearly if $\varnothing c = f(R) \cong f(R') = \varnothing d$, then $R \cong R'$, completing the proof. □

We have now shown:

Theorem 6.9. *TK is Borel complete.*

This resolves a few open questions, raised in [12]:

Corollary 6.10. *There is an ω-stable theory which is Borel complete, but which does not have ENI-DOP and is not eni-deep. Indeed it is not λ-Borel complete for any λ with $2^\lambda > \beth_2$.*

References

[1] Jon Barwise. Back and forth through infinitary logic. In Michael D. Morley and Jon Barwise, editors, *Studies in Model Theory*, volume 8 of *Studies in Mathematics*, pages 5–34. Mathematical Association of America, Buffalo, 1973.

[2] Howard Becker and Alexander S. Kechris. *The Descriptive Set Theory of Polish Group Actions*, volume 232 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, 1996.

[3] Harvey Friedman and Lee Stanley. A Borel reducibility theory for classes of countable structures. *The Journal of Symbolic Logic*, 54:894–914, 1989.

[4] S. Gao. *Invariant Descriptive Set Theory*. Chapman & Hall/CRC Pure and Applied Mathematics. CRC Press, 2008.

[5] Su Gao. *The isomorphism relation between countable models and definable equivalence relations*. PhD thesis, University of California at Los Angeles, 1998.

[6] Greg Hjorth. *Classification and Orbit Equivalence Relations*, volume 75 of *Mathematical Surveys and Monographs*. American Mathematical Society, 1999.

[7] Greg Hjorth, Alexander Kechris, and Alain Louveau. Borel equivalence relations induced by actions of the symmetric group. *Annals of Pure and Applied Logic*, 92:63–112, 1998.

[8] T. Jech. *Set Theory: The Third Millennium Edition, revised and expanded*. Springer Monographs in Mathematics. Springer Berlin Heidelberg, 2006.

[9] Alexander Kechris. *Classical Descriptive Set Theory*. Graduate Texts in Mathematics. Springer New York, 1 edition, 1995.

[10] Martin Koerwien. Comparing borel reducibility and depth of an ω-stable theory. *Notre Dame Journal of Formal Logic*, 50:365–380, 2009.

[11] Martin Koerwien. A complicated ω-stable depth 2 theory. *Journal of Symbolic Logic*, 76:47–65, 2011.

[12] Michael C. Laskowski and S. Shelah. Borel completeness of some $\aleph_0$-stable theories. *Fundamenta Mathematicae*, 2013.

[13] David Marker. *Model theory : an introduction*, volume 217 of *Graduate texts in mathematics*. Springer, New York, 2002.

[14] Robert M. Solovay. A model of set-theory in which every set of reals is Lebesgue measurable. *Annals of Mathematics*, 92:1–56, 1970.

[15] Jindrich Zapletal. *Forcing Borel reducibility invariants*. August, 2013.